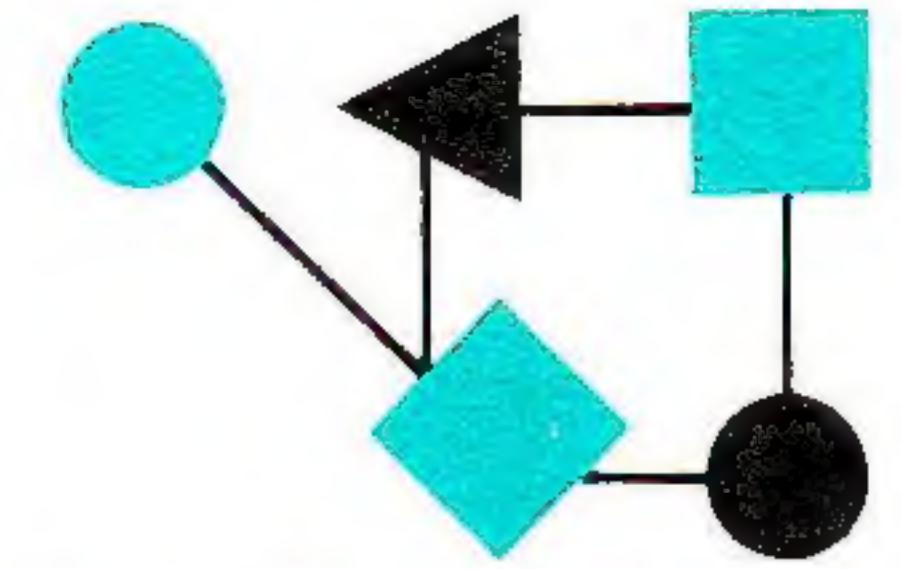


CONNEXIONS



The Interoperability Report

April 1991

Volume 5, No. 4

Connexions—

The Interoperability Report tracks current and emerging standards and technologies within the computer and communications industry.

In this issue:

Components of OSI: Systems Management.....	2
An SNMP Stereo System.....	16
International Research and Academic Networking....	20
The Ultimate File System....	28

Connexions is published monthly by Interop, Inc., 480 San Antonio Road, Suite 100, Mountain View, CA 94040, USA. 415-941-3399. Fax: 415-949-1779.

Copyright © 1991 by Interop, Inc.
Quotation with attribution encouraged.

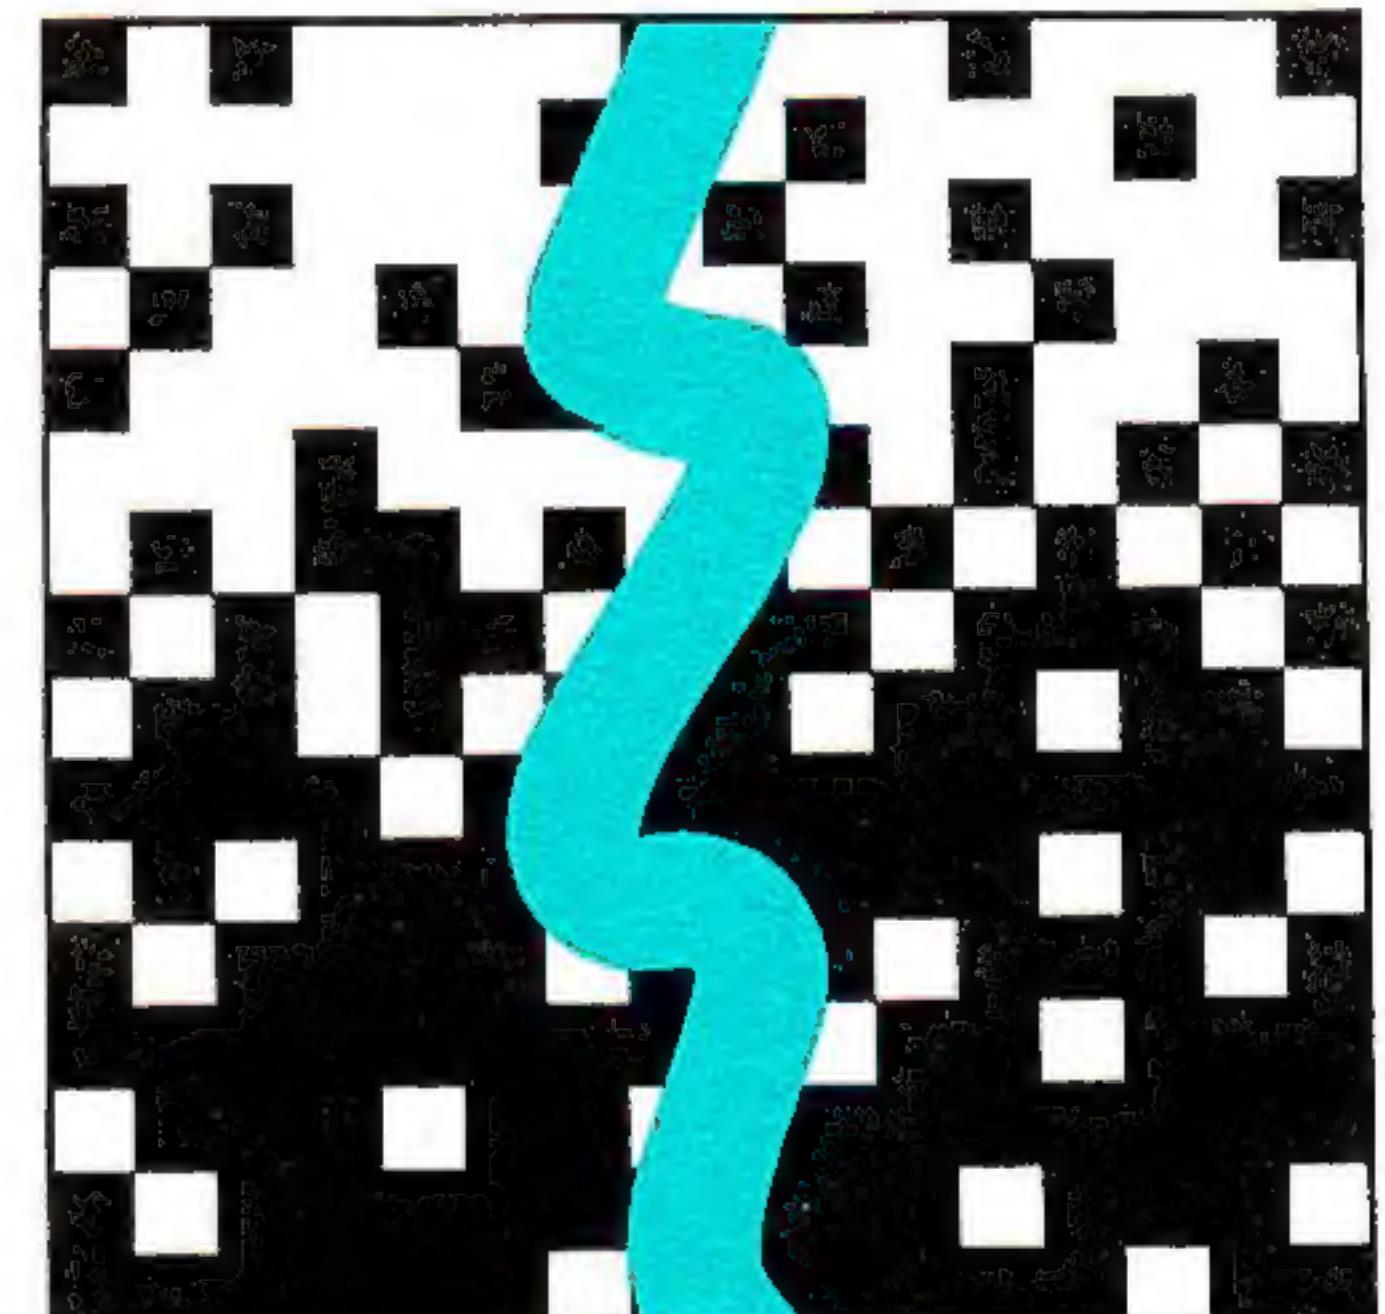
Connexions—The Interoperability Report and the *Connexions* masthead are trademarks of Interop, Inc.

ISSN 0894-5926

From the Editor

This issue of *Connexions* is being released at the *Second International Symposium on Integrated Network Management*. As a Symposium Friend, Interop Inc. is pleased to offer this complimentary copy to all symposium attendees. Several of the articles in this edition deal with network management issues. We would also like to draw your attention to two Special Issues of *Connexions*, published in March 1989 and August 1990 respectively. Both these issues focused on network management and network security, and may be obtained from our back issues department by calling 1-800-INTEROP or 415-941-3399.

Our feature article is another installment in a long-running series called *Components of OSI*. These tutorial articles have so far covered ISDN, X.400, X.500, The Transport Layer, ES-IS Routing, IS-IS Routing, The Session Service, The Presentation Layer, X.25, VT, FTAM, The Application Layer Structure, The Security Architecture, Group Communication, and more.



This month, we bring you an overview of *OSI Systems Management*, written by none other than the Symposium General Co-Chair for Advance Planning, Paul Brusil of the MITRE Corporation. It is perhaps not surprising that the management article is the largest one to date; management is both all-encompassing and complex, so one would expect the architecture for such a system to extend across all parts of the OSI Reference Model.

In the Internet Suite of Protocols, engineers have taken a somewhat different approach to network management and developed the *Simple Network Management Protocol* (SNMP). SNMP systems have been in operation for several years now, and implementors are finding new applications for this protocol as time goes on. One such application was demonstrated at our 1990 INTEROP conference and tradeshow. In this issue, Simon Hackett describes how he connected an SNMP controlled stereo system to the INTEROP Shownet.

Large internetworks are being built every day all around the world. Steve Goldstein and Christian Michau describe how academic and research networks in Europe and North America are on a convergence path.

Finally, Carl Malamud offers some thoughts on the pros and cons of comparing various kinds of file systems. His article is in response to such a comparison piece which appeared in our November 1990 issue.

Components of OSI: Systems Management

Standards at mid life; Neo Natal Products

by Paul J. Brusil, The MITRE Corporation

Introduction

Management is a critical core capability that will have major impact upon viable operational acceptance and deployment of OSI systems and networks. ISO has been developing a broad set of globally accepted management standards for management of general distributed information processing resources, including networks. Recently, CCITT, with focus on the telecom/telephony environment, and ISO, with computing industry emphasis, have agreed to progress common standards applicable to both environments. These standards have been based on object-oriented technology since their inception and are among the earliest international standards using this technology.

The broad mix of standards from increasingly diverse communities is making OSI systems management the world wide basis for multi-vendor, interoperable, integrated management. With interest growing in organizations like OSF and X-Open, the span of OSI integrated management will widen from today's integrated management of network elements in public (carrier-provided) and private LANs and WANs to include system and application software.

A large number of these standards are now stable and mature. Some have taken the volume of standards to mean high implementation costs. Others will leverage rapidly increasing memory availability and CPU power, as well as falling costs, to implement these standards in traditionally resource-constrained bridges and modems.

Early adopters

The commercialization of these standards has begun. The first multi-vendor demonstration of early OSI management standards occurred at INTEROP 88 (see Marshall in [2]). Numerous Japanese firms including NEC, Fujitsu, Hitachi, Toshiba, IBM-Japan and several others are now completing efforts to demonstrate mature OSI management and advanced OSI protocols such as *Transaction Processing* and *Remote Database Access*. OSI management products can be bought now; some were available for purchase at *ComNet '91*. Many industry giants will offer products soon. Those leading the commercialization charge include IBM, Digital, AT&T, British Telecom, Hewlett-Packard, Unisys, NCR, as well as the other 100+ international vendors in the OSI/Network Management Forum and the other 150+ international organizations within the *Network Management SIG* (NMSIG) of the Department of Commerce's *OSI Implementors Workshop*. Vendors and testing firms in nearly 20 countries are committed to OSI management as the method for enabling networks to access and manipulate management information on a worldwide basis.

This article provides the basic concepts associated with OSI systems management. It gives a tour through the various associated CCITT and ISO joint standards. It also considers work of organizations contributing to commercialization of this technology.

Basic System Management concepts

General concepts of OSI management have been described in several contexts, e.g., [1, 2], and are summarized in Figure 1. Basic components of management include *manager* and *managed* (often called *agent*) systems. These systems may contain the actual resources that provide the "useful work" in the distributed system being managed. For example, protocols provide the useful work in network elements. Sometimes these systems are management *proxies* that act as intermediaries between actual resources and management systems.

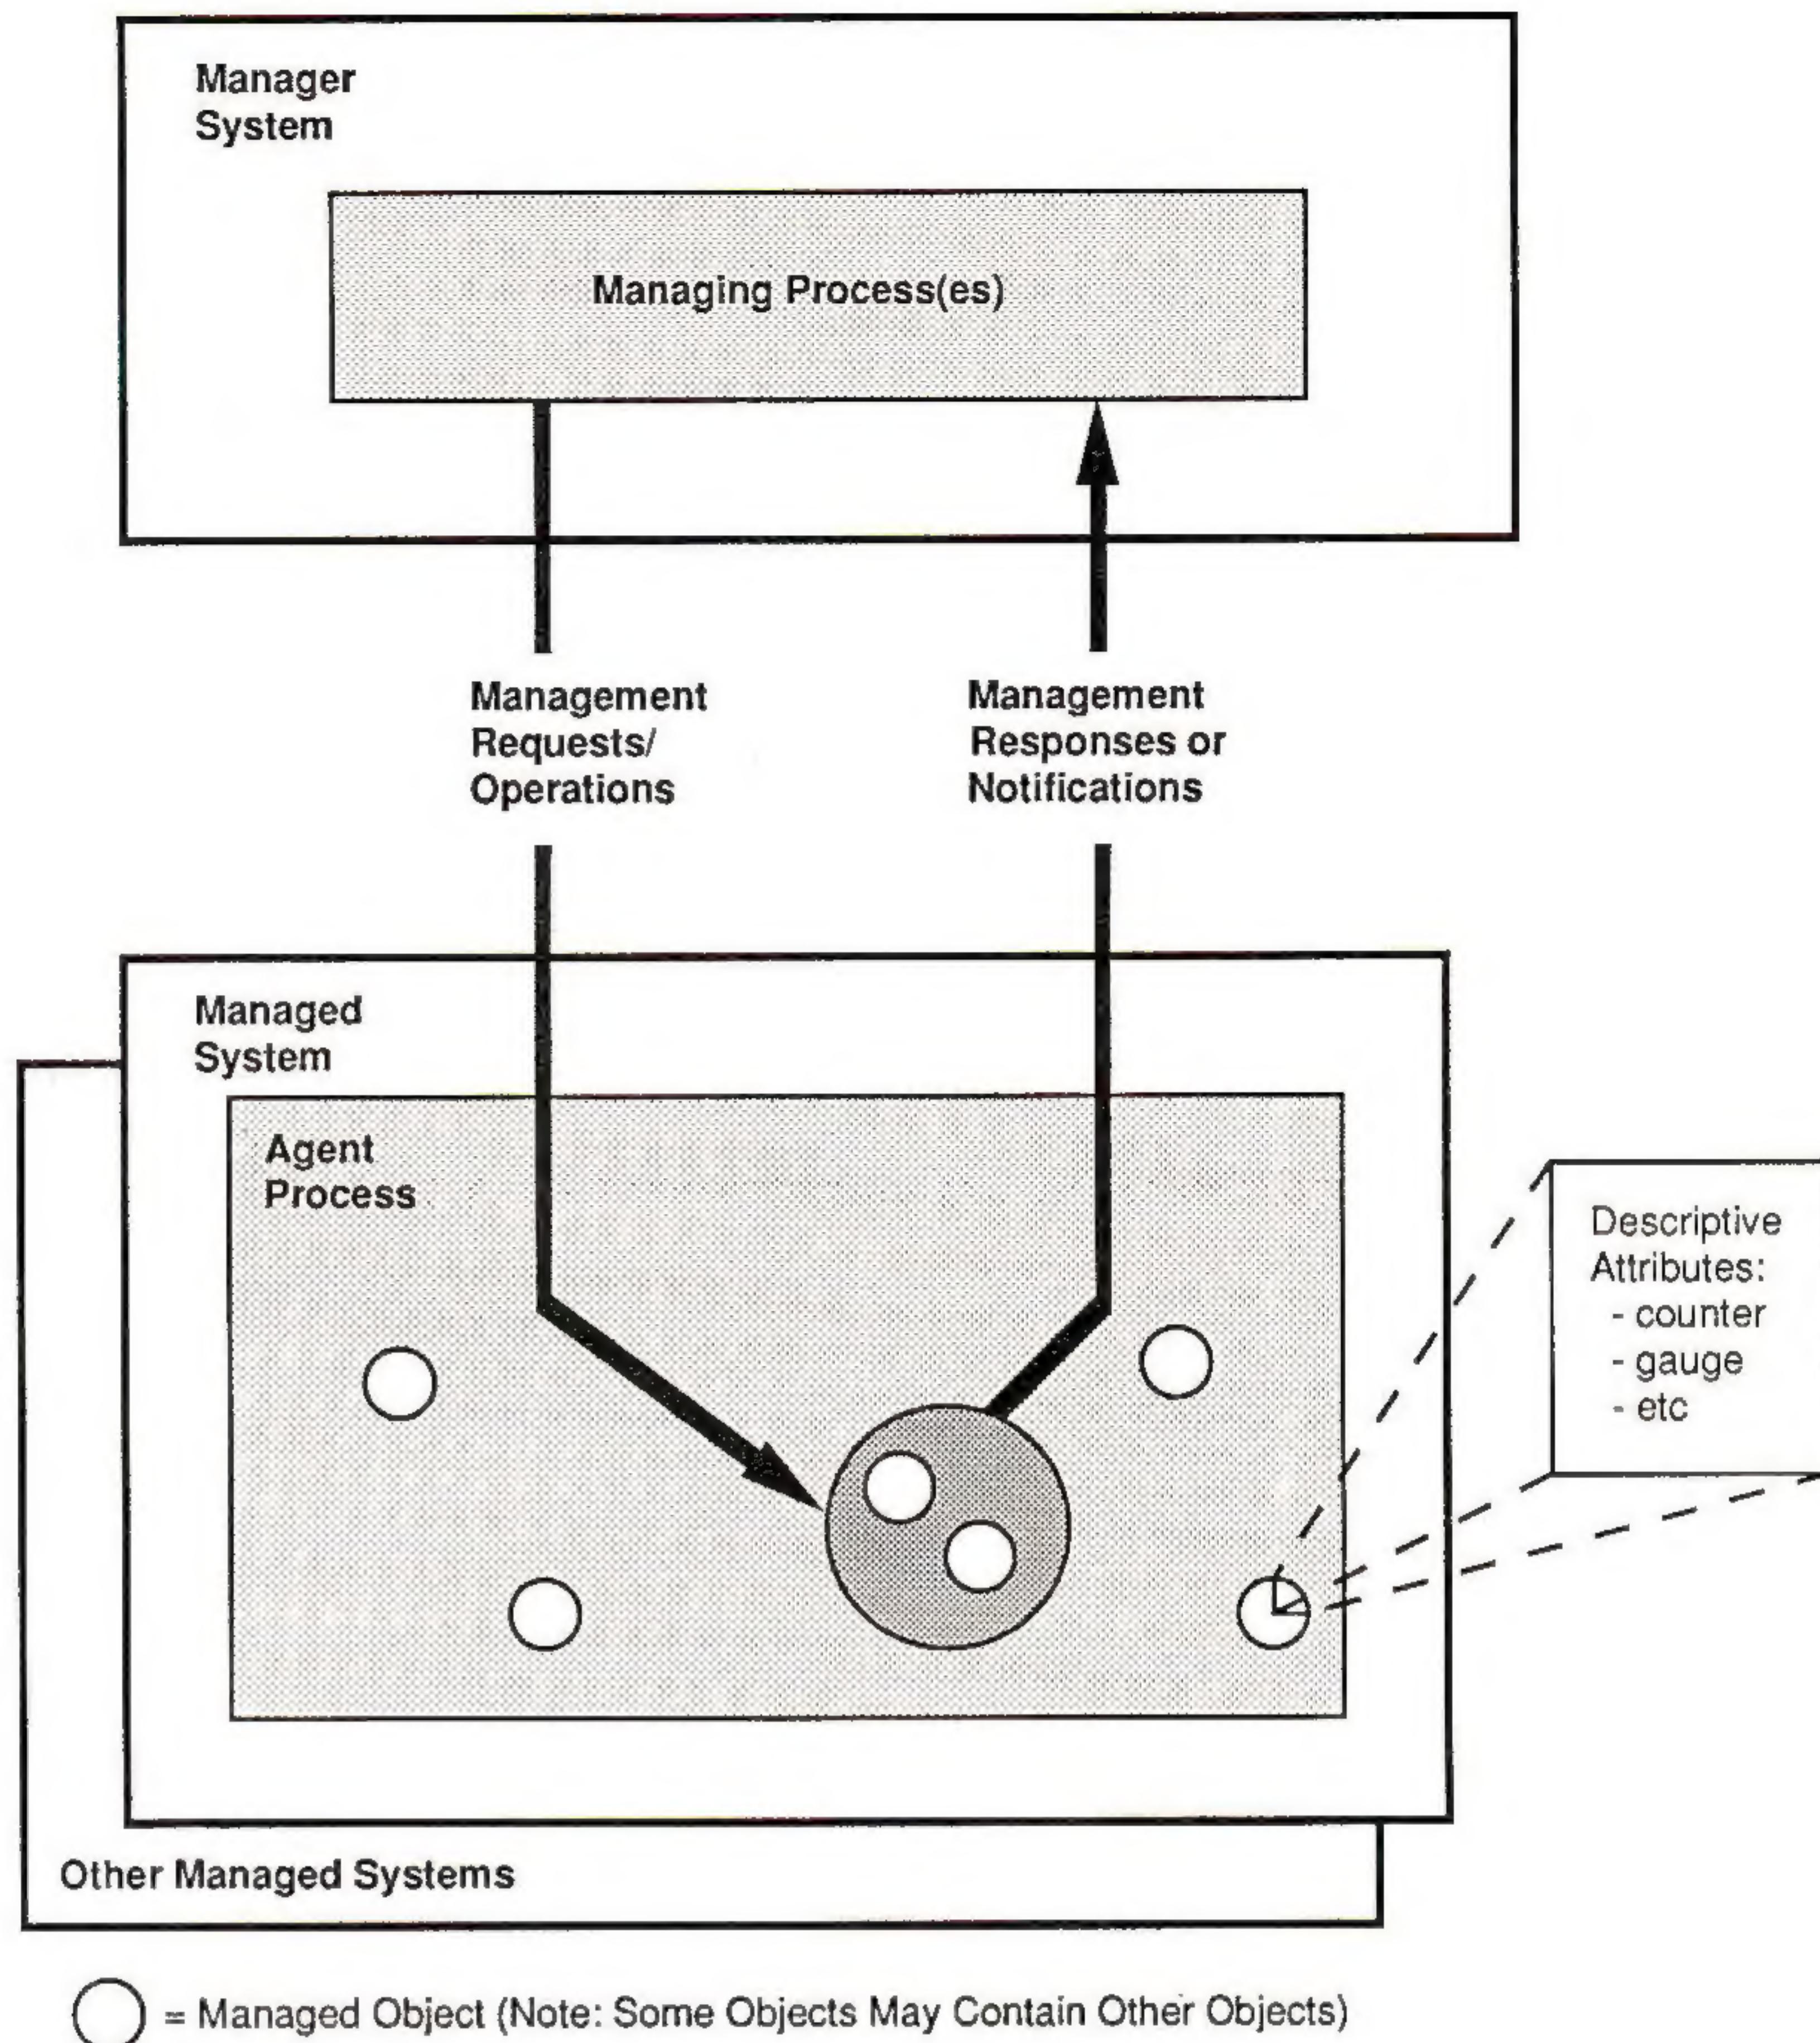


Figure 1: OSI Systems Management Framework

Managed Objects

Management activities may occur via application layer, manager processes in manager systems interacting with remote managed objects in managed systems. Through such interactions, one manipulates or retrieves management information pertinent to the resources being managed. The interactions occur via remote agent processes in managed systems. According to object-oriented paradigms, resources are modeled as *Managed Objects*. Managed objects logically characterize the actual resources for purposes of managing the resources. That is, a managed object models the permissible management operations (e.g., read, replace) on the managed object, permissible notifications that the object may wish to send, and descriptive attributes (e.g., counter, gauge) associated with the managed object. Managed objects may contain other managed objects.

For meaningful interoperation, manager and agent processes must share the same schema of knowledge about the managed objects in which they are mutually interested. Different schemata are possible so that different manager processes can have different management views of the same managed object.

Agent processes perform the management operations requested by manager processes, e.g., reading attributes of a specific object, setting object attributes, or returning a response to the corresponding requesting manager process. They may forward, according to specifiable rules, certain notifications (events) asynchronously generated by their associated managed objects. Agent processes may interact with their associated managed objects in ways not subject to standardization since such interactions may take place completely within one implementation.

continued on next page

Example: Managing networks

OSI Systems Management (continued)

For a system of layered communication protocols, the OSI Systems Management framework recognizes three management approaches: management within a communications protocol, management within a layer (with several protocol instances) and management across layers. Protocol management, such as transport flow control windowing, consists of protocol-internal mechanisms needed to control a particular instance of communication. Layer management consists of layer-specific activities and networking services needed to manage resources associated with a particular layer, e.g., network layer routing mechanisms or link layer token control. Layer management acts directly at a single layer and does not necessarily rely on application layer manager and agent processes. Systems management pertains to management of networking resources associated with several protocol layers and protocol instances. All approaches may occur concurrently in the management of network systems.

Figure 2 conceptualizes OSI management for network elements. Manager and managed system elements contain the stack of communication *Protocol Entities* (PE) that provide networking services, as well as managing and agent application processes, respectively. These processes can access local managed objects and can communicate with remote partner processes to access remote managed objects. The PEs are some of the distributed resources to be managed. Associated with PEs may be *Layer Managers* (LMs) and management information. Managed objects may be layer objects specific to PEs; or, they may be system objects pertinent to several layers of PEs and/or the network element as a whole. The conceptual aggregate of all such management information is called a *Management Information Base* (MIB).

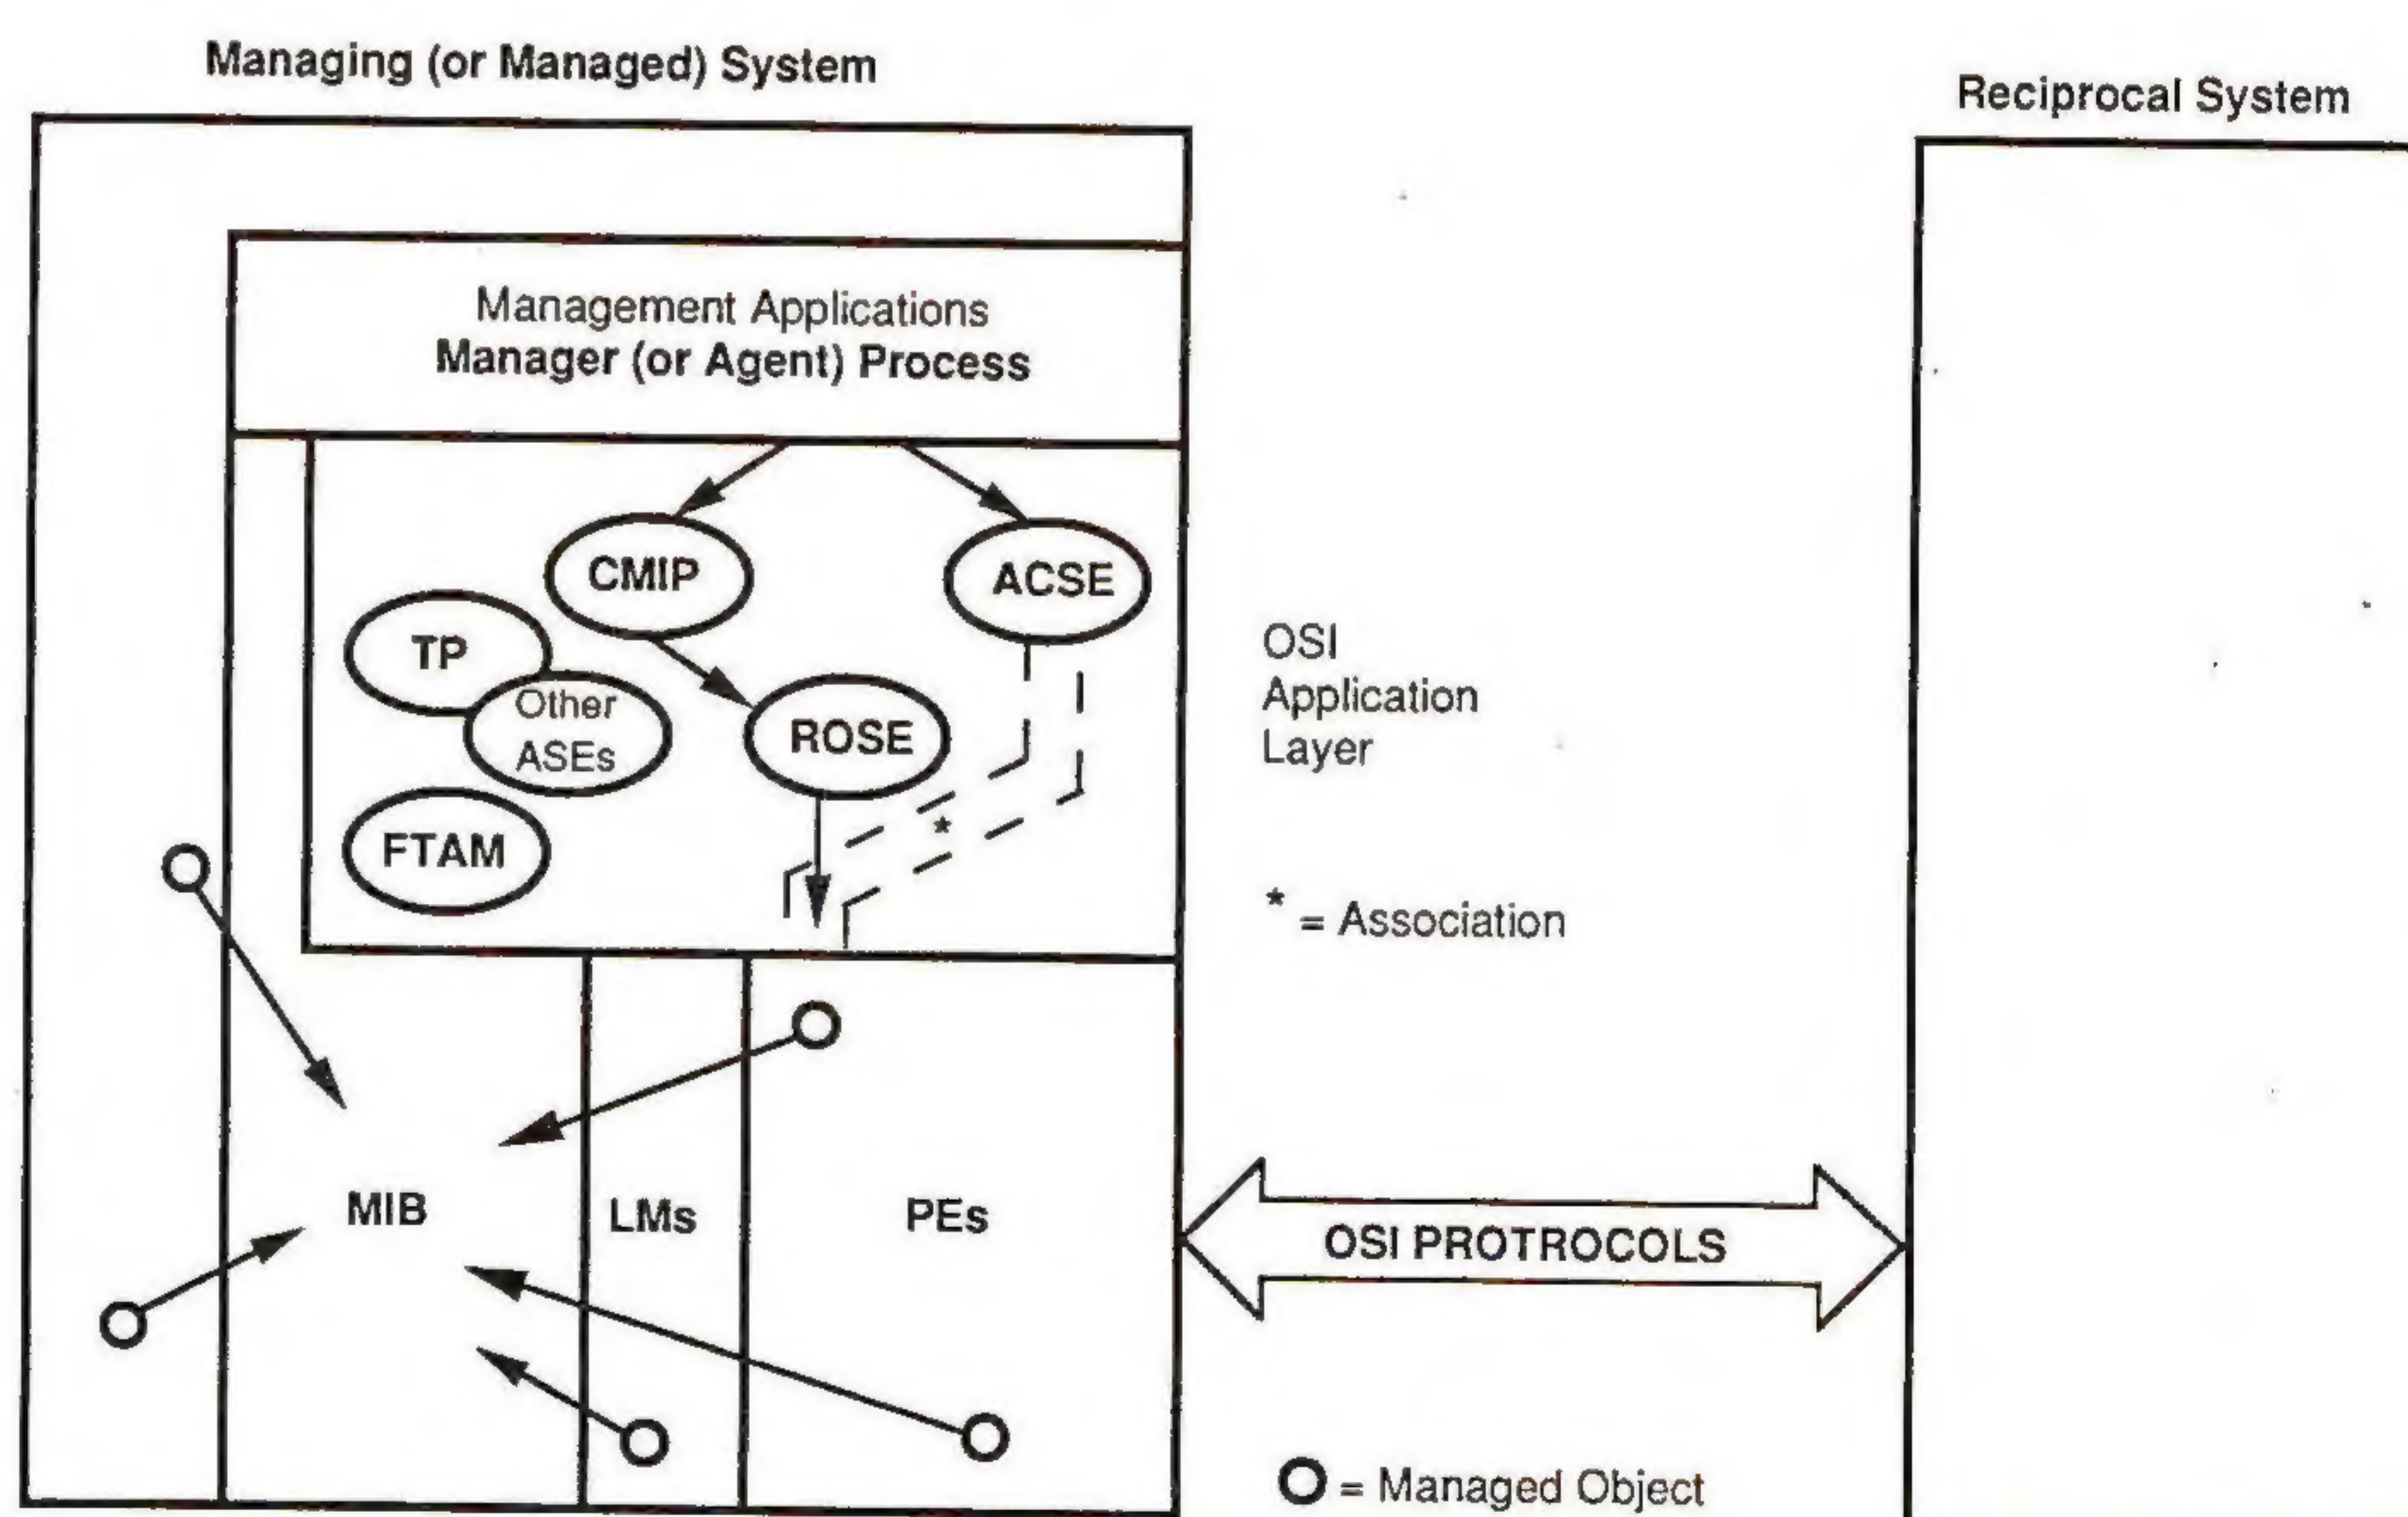


Figure 2: Architectural Model of OSI Management in Manager and Managed System Network Elements

Managing and agent processes rely on numerous OSI *Application Service Elements* (ASEs) which work together, calling upon each other as described in [3], to effect management communications. Typical of OSI, communications are connection oriented. An *Association Control Service Element* (ACSE) establishes and terminates associations (application-to-application “connections”) between the managing and agent processes.

Management service requests and responses between the processes are then effected by the *Common Management Information Protocol*, CMIP, whose packet data units are sent over such connections. A *Remote Operations Service Element* (ROSE) provides a transaction-oriented, request/response service (similar to SNMP's Top-level Message) to carry the CMIP operations over the association. ROSE also provides for linking multiple replies to a single request. It has a parameter for indicating which CMIS operation is carried in the Remote Operations request or response.

Managing and agent processes may use services of other ASEs. For example, a transaction processing ASE could provide coordination and concurrency control of distributed management transactions if managers choose to synchronize changes across distributed resources. A file transfer ASE could provide bulk transfers for software configuration loads or transmission of aggregated management reports.

The context of which set of ASEs can be used by remote cooperating management applications is established at the time that an association is set up between the applications.

Operationally, the manager process may receive inputs from local administrative personnel or their software agents, from local LMs, from remote agent processes, and/or from remote LMs. Decisions made by the manager process are either effected by local mechanisms to the local managed objects, or they are communicated to remote agent processes or LMs.

Standards overview

The process of developing OSI standards has swung from codification of de jure, industry-wide approaches to international amalgamation of the autonomous research results of standards' participants. Furthermore, OSI System Management standards are being produced by many national and international standards organizations, including ANSI, several ISO sub committees and working groups, several CCITT study groups, and the IEEE. (Chappell [4] overviews relationships among these organizations.) Accordingly, the various standards are progressing at different rates, often perceived as slow. Some justify a deliberate pace because of the broad scope, complexity and intricacy of the subject, which cuts across all protocols, layers, numerous devices, and various technological disciplines, organizations, countries and laws. Others view the pace as arising from the difficult organizational problems and language barriers attendant in creating worldwide standards. ISO standards themselves represent a compromise among 87 countries representing more than 95% of the world's industrial production. Open consensus-building in such diverse forums is truly difficult. Accordingly, OSI standards in general, and systems management standards in particular, are often as much political achievements as they are technical compromises.

OSI Systems Management has four groups of standards: (a) those summarizing basic management concepts and models as well as the other standards that further specify details, (b) those relating to the specification of managed objects, (c) those specifying systems management functions that provide value-added capabilities beyond basic management communication services, and (d) those pertaining to the basic application layer services and protocols for communicating management information. In brief, these standards specify managed objects and attributes, as well as rules for defining, manipulating, sending information about, and controlling the sending of information about managed objects and attributes.

OSI Systems Management (*continued*)

Within each group there are several standards. Currently, over 100 international standards are being developed. ISO is developing over 30 standards pertaining to OSI Systems Management; CCITT is developing over 60, 23 of which are identical to corresponding ISO standards; IEEE 802 and P1003 are developing at least 7. Figure 3 shows the standards being developed within one ISO *Sub Committee* (SC). They form the core set upon which most of the other standards are based. Their contents, inter-relationships and inter-dependencies are summarized in the next section.

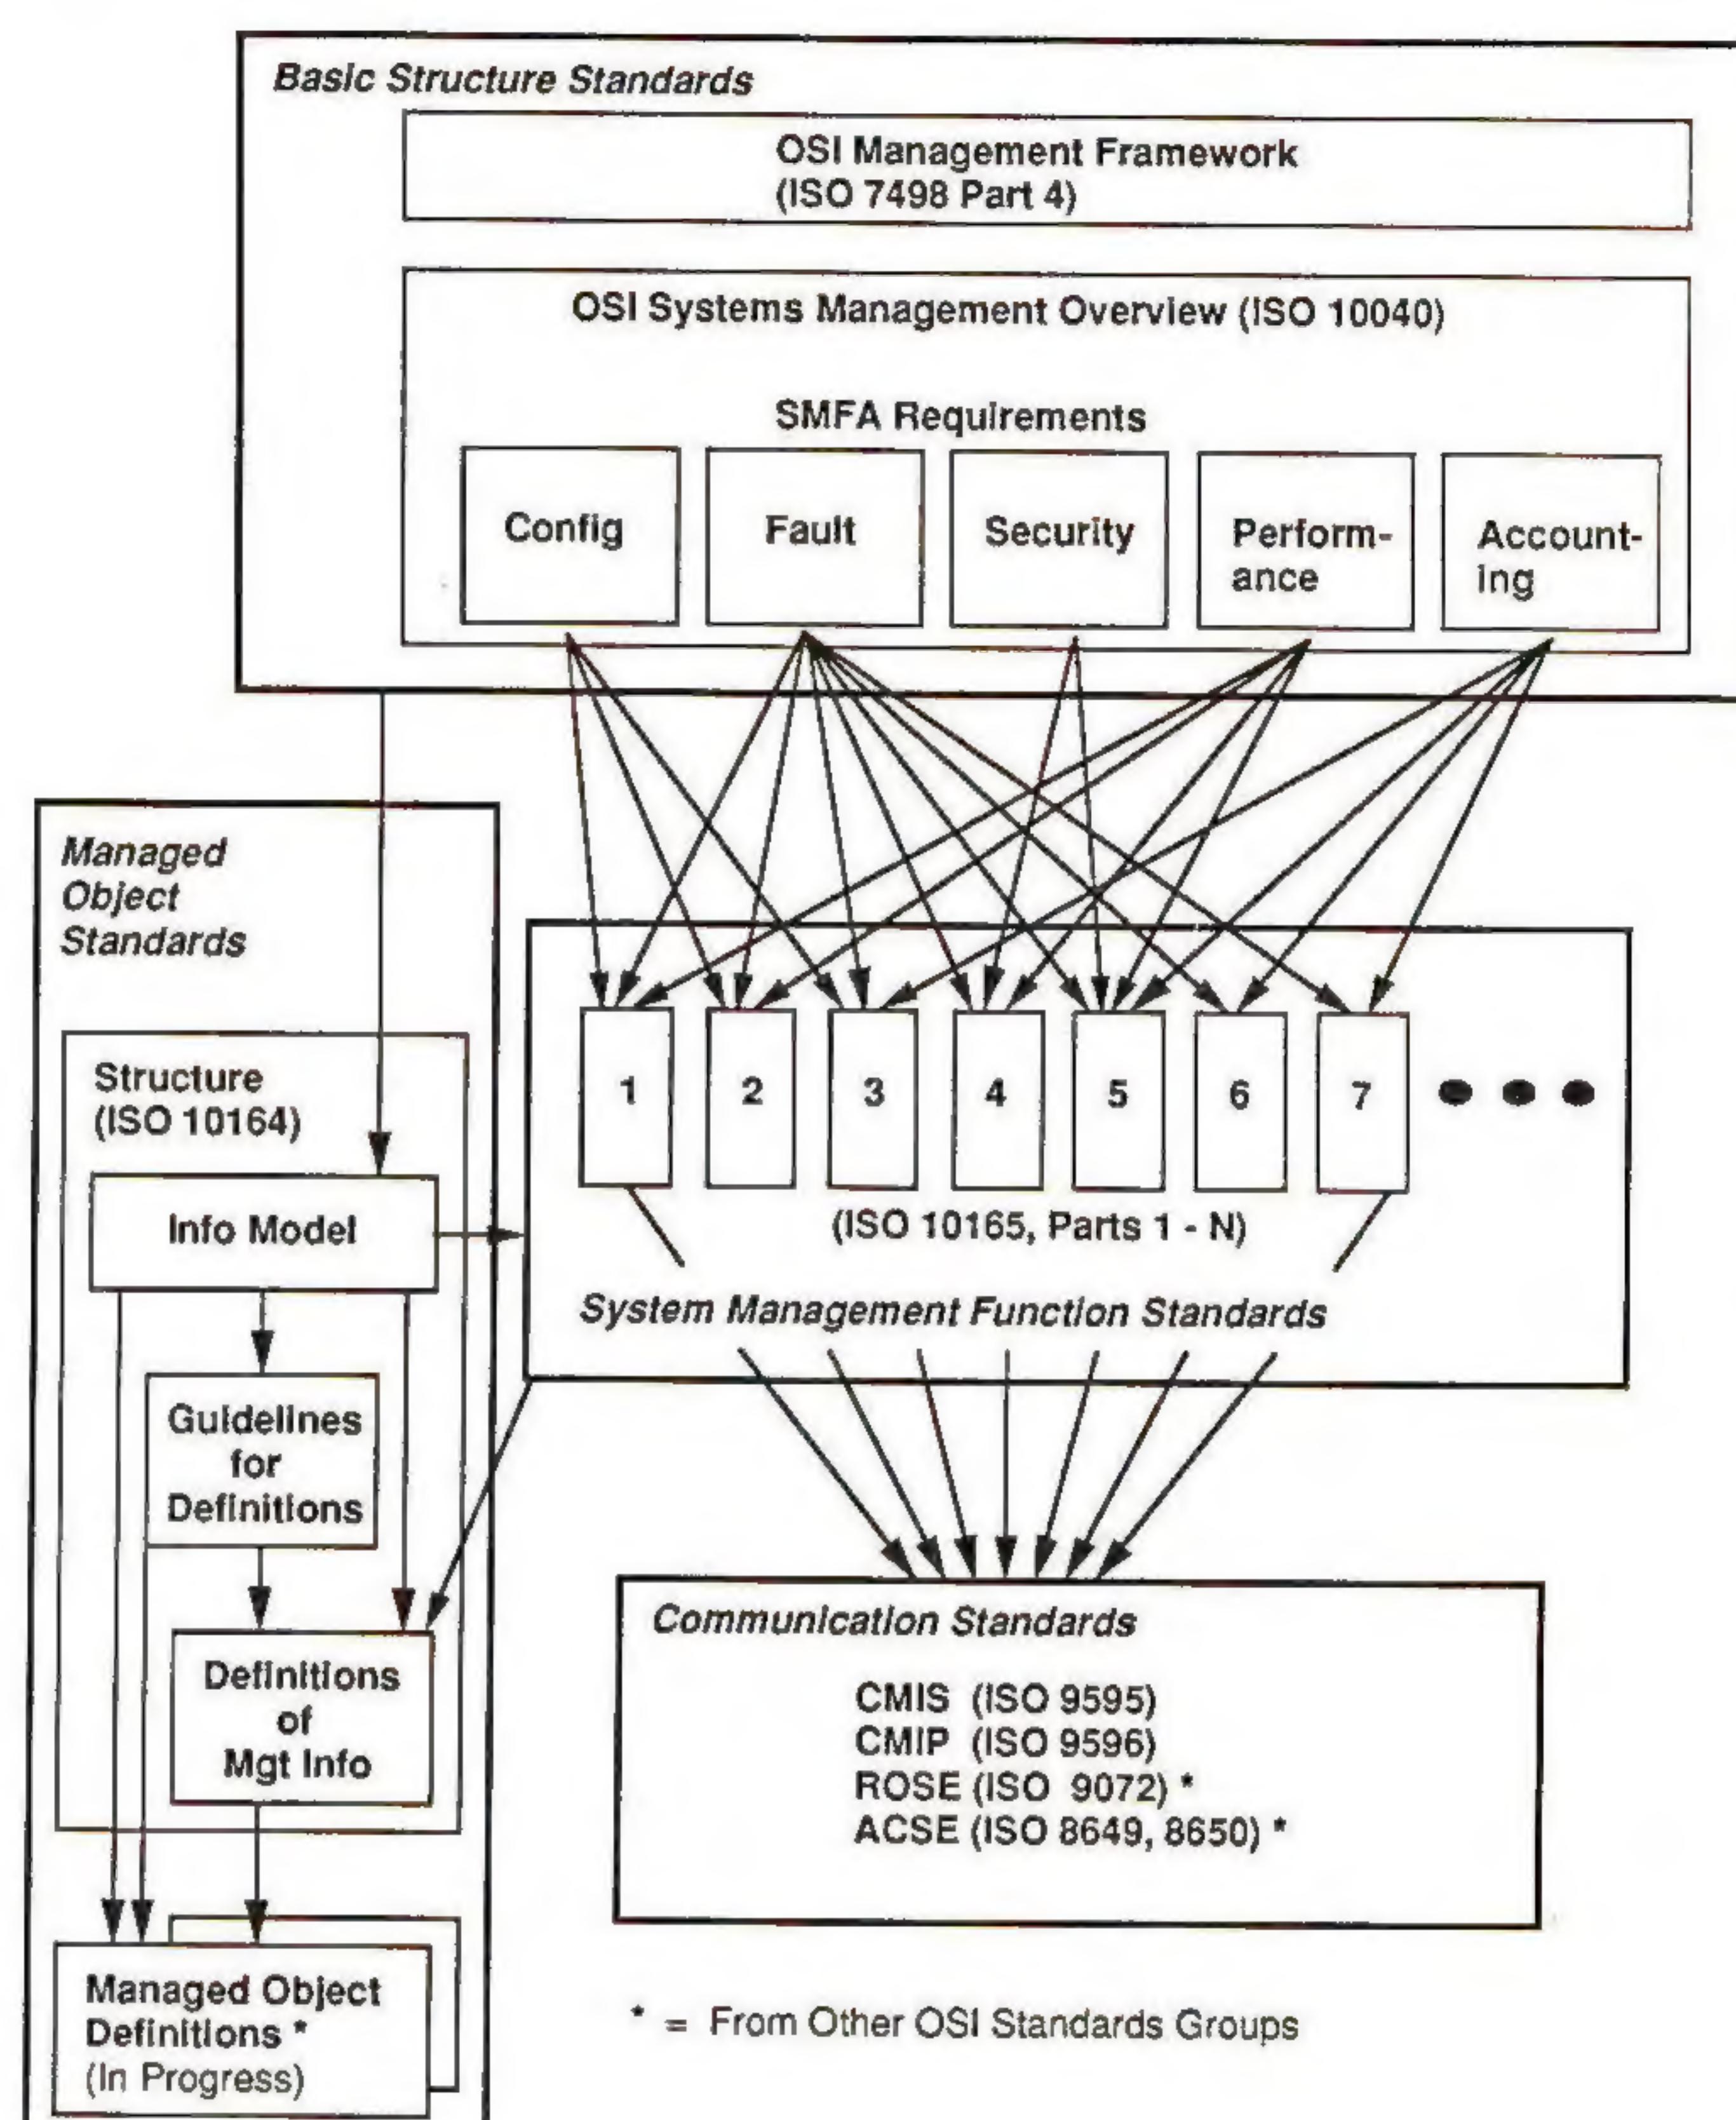


Figure 3: Core OSI System Management Standards (From ISO SC21)

Good news: there are now 18 stable ISO international systems management standards at the *Draft International Standard* (DIS) or *International Standard* (IS) levels. All 11 DISs are scheduled to reach IS by mid 1991. These 18 standards provide such a sufficiently rich and stable set of management tools that implementations of OSI management systems can safely begin without fear that they may be obviated due to changes in evolving standards.

A number of the less mature standards, *Committee Drafts* (CDs), are planned to be technically stable enough to become DISs in 1991. Today's major short-coming for OSI management is that stable, internationally-standard definitions of many managed objects, such as those progressing in CCITT, IEEE and ISO, are not yet in hand.

Guide to standards

Details within the core standards are too voluminous to summarize here. Only salient features and a guide to the standards are given here. Indicated with each standard's description are the ISO document number, the current standardization level, the anticipated IS date, and the corresponding CCITT recommendation number.

Management Structure standards

These documents are available from the *American National Standards Institute* (ANSI) in New York City, or *Omnicom, Inc.* in Vienna, Virginia.

Management Framework (ISO IS 7498-4, OSI Basic Reference Model; CCITT X.700). This and ISO standard 10040 provide the architectural information required to guide development of all the system management standards. The framework defines major elements of OSI management: *System Management Functional Areas* (SMFAs), MIB, basic managed object concepts, and different methods of management information interaction, i.e., system, layer and protocol management.

SMFAs enumerate generic requirements for management in several areas (see, e.g., Stine et al in [1]). For example, requirements to enable identification or negotiation of mechanisms for associating and collecting communication resource usage charges, to initiate or deactivate charging algorithms, and to monitor or report account relevant information are given in an *Accounting Management SMFA*. The *Security Management SMFA* requirements include capabilities to combat threats by providing audit trails of security-relevant events, capabilities to analyze such audit trails, and tools for supporting the control of security services and mechanisms as well as the associated decision making, e.g., when to redistribute keys, when to reinitialize encipherment algorithms and when to isolate infected nodes. Other SMFAs—*Fault Management, Performance Management and Configuration Management*—contain obvious requirements to detect and to isolate problems, to monitor and to tune networks, and to monitor and to maintain continuous operations, respectively.

Systems Management Overview (ISO DIS 10040, IS in '91; CCITT X.701). This standard expands the framework's concepts via a systems management model. The model refines managed objects concepts and provides a brief guide on how to use the management-information-related standards to define managed objects. The model addresses functional aspects of systems management by describing the relationships between SMFA requirements and the system management functions standards (ISO 10164). It addresses communication aspects such as management processes acting in either a managing or an agent role. It also addresses application layer concepts relating to systems management and the establishment of shared knowledge about the intrinsically distributed elements of a management system. Lastly, it addresses organizational concepts, such as management domains, that describe ways management can be distributed for reasons of management scale, security, accounting, or the need to provide administrative autonomy over an organization's resources. This standard also provides a brief overview of the other systems management standards.

Management Information standards

Standards defining management information pertain to the *Structure of Management Information* (SMI) and the managed object descriptions applicable to real resources. Rules guiding those developing the definition, structure and identification of managed objects are specified in the former (ISO 10165)—a multi-part standard. Having objects defined according to a common structure is perhaps the single most important notion that will facilitate integration of management across a variety of heterogeneous resources and system environments. This notion transcends the significance of a single common management communication protocol within such heterogeneous environments.

OSI Systems Management (*continued*)

SMI—Part 1: Management Information Model (ISO DIS 10165-1, IS in '91; CCITT X.720). This standard defines (a) a model governing logical organization of management information and relationships into which managed objects can enter (inheritance, specialization, allomorphism, containment), (b) consistent and versatile naming and containment principles so that managed objects and their attributes can be identified in, and accessed by, management protocols, and (c) so-called Managed Object Knowledge tools to allow dynamic discovery of which managed objects have been instantiated in a given system.

The information model is heavily based on object-oriented design principles. Accordingly, it permits easy extensibility to incorporate new classes of managed objects and functions as they emerge, modularity and extensibility of management protocol and procedures, and reuse of pieces of management information specifications.

Attributes, the properties of managed objects, may be either *single-valued*, as in some of today's other management approaches, or *set-valued*. Set-valued attributes are those whose value is a set of members of a given datatype. For retrieval purposes a group of attributes may be referenced by a single "group" attribute. Attributes determine or reflect behavior of managed objects. They may be collected together with notifications, operations and behavior to form conditional packages that can be conditionally instantiated in managed objects.

Objects can be extensions of other objects by inheriting characteristics, such as operations, attributes, behavior, from superclass objects, and by adding new characteristics such as attributes, actions, action arguments, and/or extending/restricting, e.g., attribute and argument ranges. Allomorphism, a technique to extend object classes to support new equipment and technology types without making older management systems obsolete, is specified so that migration between versions of management is possible.

The model allows two types of management operations. The first can be applied to object attributes (get/replace/set-to-default attribute value and add/remove members of set-valued attributes). The second apply to whole objects (create/delete/action).

Lastly, principles of containment and naming are specified to identify managed objects uniquely. Managed object identifiers are of two components: one identifying the class of objects to which a managed object belongs (e.g., a transport entity object class defined by an ISO standards body), the other identifying the specific instance of the managed object (e.g., the transport object in a specific network element). The former is an ASN.1 object identifier, a sequence of integers that navigate through a hierarchical registration tree to the class of the managed object (see, McCloghrie et al in [2]). The latter is a distinguished name based on the concatenated sequence of relative distinguished names of object instances in a containment tree (see, e.g., [7]).

SMI—Part 4: Guidelines for the Definition of Managed Objects (ISO DIS 10165-4, IS in '91; CCITT X.722). This standard, generally referenced by its acronym GDMO, specifies a set of guidelines and templates to be used by individuals defining managed object classes. It, therefore, encourages consistency among object definitions, ensures definitions compatible with all appropriate OSI management standards, and reduces duplication of effort among object definers by identifying commonly useful documentation layouts, procedures and definitions.

By using these templates and the grammar and rules of the ASN.1 specification language, managed object definitions will consist of enumeration of:

- the object's attributes,
- the operations that can be applied to the object,
- any conditions/constraints that may affect execution of each operation (e.g., any synchronization constraints within the object and criteria for supporting operation requests for atomic synchronization with other objects),
- the effects (exhibited behavior) that these management operations have upon the object and its attributes as well as upon related objects,
- the notifications that can be emitted by the object,
- the behavior and syntax associated with specific action type of operations,
- the conditional packages which can be present in the object,
- the position of the object class in the inheritance hierarchy,
- alternative naming structures/bindings that define possible naming relationships with superior objects and the managed object classes from which subordinate objects may be instantiated, and;
- the objects that are allomorphic with this object class.

This standard also summarizes requisite registration concepts for object definers. Contrary to other management approaches, OSI managed object naming and registration are independent.

SMI—Part 2: Definition of Management Information (ISO DIS 10165-2, IS in '91; CCITT X.721). This standard defines generic attribute types, specific attributes, and packages, including their syntax, behavior, valid operations, as well as name bindings for object classes. These definitions are specified via GDMO templates.

This standard enumerates a set of support managed object classes, most of which are fully specified in the system management function standards (ISO 10164). These support objects, attributes and packages may in principle be reusable components within definitions of other managed objects in a wide variety of classes. Examples of such support management information definitions include:

- a discriminator object for controlling the filtering of notifications that may be reported as events to a remote manager or log,
- generic attribute types such as counter, gauge, counter threshold, gauge threshold and timemarks, and
- generic packages for daily or weekly scheduling of, e.g., notification filtering.

Managed Object definitions

In addition to support managed objects discussed above, every standards group that specifies a protocol or a component of a communications system is responsible for specifying, according to SMI rules, the managed object(s) and attributes associated with that resource. The number of such standards groups is very large.

OSI Systems Management (*continued*)

Those actively engaged in defining objects according to the SMI/GDMO include ISO's transport and network layer bodies, ANSI's FDDI committee, CCITT's and ANSI's telecommunication management network bodies, IEEE's 802 group (work is underway in 802.2, 802.3, 802.5, 802.6 and 802.10) and IEEE's P1003 POSIX group.

Many industry consortia are also using SMI/GDMO to define system objects required for multi-vendor interoperability, as well as experimental, pre-standards, layer objects. Active groups include *INTAP* (Interoperable Technology Association for Information Processing—Japan) defining objects for ISDN and FDDI internets, *CNMA* (Communications Network for Manufacturing Applications) defining OSI upper and lower layer objects for European experimentation, the *OSI/Network Management Forum* defining objects for manager-to-manager interactions, and NMSIG defining objects for high level management of internets with heterogeneous lower layer protocols.

To encourage early and quick, interim managed object implementations, some in ISO are considering to define generic managed objects. Comparable to ongoing NMSIG initiatives, these would be high level overview definitions of key system and networking resources. They would be developed before the detailed standard managed objects are internationally hammered out by the various resource-specific standards bodies. Early management implementations could include these generic managed objects. By themselves, they could accommodate high level, "surface" management. Additional interim refinements by vendors could provide more meaningful objects incorporating resource-specific details. Only software pertinent to the interim refined definitions would need to be swapped out when standard refined definitions become available. This would permit early fielding of objects without significantly hampering commercial transition to subsequent standard objects.

IMIL

To encourage open and non-redundant definitions and refinements of these and/or other vendor-specific objects, the NMSIG has been working with its regional sister organizations, the OSI Workshops in Europe and Asia, to develop notions of an *International Management Information Library* (IMIL). The IMIL is to be a widely and readily accessible repository that lists or points to all object definitions in the standard GDMO syntax. Standards governing establishment and maintenance of the IMIL would need to be developed.

Near term uses of the IMIL would be many. It would give wide spread public disclosure to managed objects of all sorts, including those to manage multi-vendor interoperable applications, such as spread sheets and graphics, and multi-vendor interoperable system software such, as data base management systems. It would be used to reduce proliferation of comparable object definitions by screening for, and potentially working to harmonize, such similarities before definitions are entered into, or referenced by, the IMIL. In true object oriented fashion, the IMIL would promote reuse and refinement of existing management information definitions. Once the IMIL becomes richly populated, vendors will be able to manage both standard and, at least to some extent, proprietary resources of other vendors. Additionally, users will have a single complete directory and dictionary of OSI manageable objects from which procurement selections can be made. There is already discussion about using profiles of IMIL objects as part of emerging US government federal information processing standards to be mandated soon for government management procurements.

System Management Functions standards

The multi-part, *System Management Functions* (SMF) standards define functionality to support SMFA requirements. A SMF may support requirements in many SMFAs, e.g., the *Event Report Management Function SMF* may be applicable to all SMFAs. Conversely, a SMFA may require several SMFs. Each SMF standard provides mappings between services provided by the SMF and CMIS. Initial SMFs being progressed are as follows.

Object Management Function (ISO DIS 10164-1, IS in '91; CCITT X.730). This standard specifies services for creating, deleting, examining and changing sets of management information. It specifies services for reporting creation/deletion of managed objects, name changes to managed objects, and changes to attribute values.

State Management Function (ISO DIS 10164-2, IS in '91; CCITT X.731). This standard models and specifies the states which managed objects may have. It defines nine generic attribute types useful for inquiring about and changing the management state of a managed object. One notification (event) type is defined for reporting management state changes when they occur either through internal operation of, or management action upon, the resource. Operational, usage and administrative viewpoints are addressed. Services are defined for monitoring operability and usage of system resources, for administratively restricting their availability, and for restricting receipt of state change notifications.

Attributes for Representing Relationships (ISO DIS 10164-3, IS in '91; CCITT X.732). This standard models and identifies types of relationships which can exist among managed objects representing different parts of a system. It specifies ten generic attribute types and one notification type, together with their applicable parameters and semantics for importation into managed object definitions when relationships with other managed objects are to be specified. Services are defined for establishing, examining and monitoring the relationships among managed objects, and therefore for observing how operation of one part of a system depends on other parts.

Alarm Reporting Function (ISO DIS 10164-4, IS in '91; CCITT X.733). This standard models alarm reporting. It specifies five generic alarm notifications (events), together with their parameters and semantics. These notifications are associated primarily with fault management.

Event Report Management Function (ISO DIS 10164-5, IS in '91; CCITT X.734). This standard provides a model for event reporting and the control of event reporting. It specifies means for controlling selection and distribution of events to manager-specifiable destinations. It specifies an event forwarding discriminator managed object that defines manager-creatable/setable criteria by which managed object notifications may be conveyed remotely as event reports, as well as the time periods during which such event forwarding discrimination can occur. LaBarre [5] gives an excellent tutorial. Event-driven management is deemed particularly crucial in WANs with premium transmission costs and in bandwidth-constrained aeronautical and military environments.

Log Control Function (ISO DIS 10164-6, IS in '91; CCITT X.735). This standard provides a model for logging events and other management information. It specifies a log, and services by which records associated with event reports can be logged.

OSI Systems Management (*continued*)

Similar to concepts associated with event forwarding discrimination, records can be logged according to numerous manager-setable schedules and only if manager-creatable/setable logging criteria are passed. See [5].

Security Alarm Reporting Function (ISO DIS 10164-7, IS in '91; CCITT X.736). This standard models reporting of security-related events (e.g., attacks on, or breaches of, system security) and misoperations in security services and mechanisms. It specifies generic security alarm notifications, together with their parameters and semantics, as well as facilities for creating, deleting and modifying event forwarding discriminators for controlling selection and distribution of security alarms to manager-specifiable destinations.

Security Audit Trail Function (ISO CD 10164-8, IS in '92; CCITT X.740). This standard provides extension of the log control function standard, with its discriminator concepts, to security relevant event logging. It specifies how to control starting/stopping of security auditing and creating/modifying auditing criteria. It specifies generic security audit trail notifications and their applicable parameters and semantics.

Objects and Attributes for Access Control (ISO CD 10164-9, IS in '93; CCITT X.741). This standard models access control for management communication associations, as a whole, as well as for individual management operations within an association. It specifies managed objects and attributes (to be associated with those objects to be managed and protected in a system) to be used to grant or to deny access according to the access control policy represented by this access control management information.

Accounting Meter Function (ISO CD 10164-10, IS in '92; CCITT X.742). This standard provides a model for accounting meters and logs, and for the control of such objects. It specifies services for retrieving, reporting, and recording resource usage data and for selecting which usage data are to be collected and under what conditions they are to be reported.

Workload Monitoring Function (ISO CD 10164-11, IS in '92; CCITT X.739). This standard models gauges to be used for resource utilization monitoring, for rejection rate monitoring and for resource request rate monitoring. It specifies gauge managed objects, a mean monitor metric object for deriving the instantaneous mean value of an associated gauge object, the notifications that can be emitted, and operations for initiating, terminating, suspending, resuming and modifying metric monitoring.

Test Management Function (ISO CD 10164-12, IS in '92; CCITT X.745). This standard supports remote control of tests involving open systems and the specification of tests which exercise OSI resources in such systems. Individual tests may be used in the context of several different SMFAs, such as both fault and performance management.

Measurement Summarization Function (ISO CD 10164-13, IS in '92; CCITT X.738). This standard provides a model for sampling and aggregating (in manager-specifiable ways), optionally buffering, and reporting (according to manager-specifiable schedules) various summarizations of values of manager-specifiable attributes of (a) specific object instances across time (time averages), or (b) a set of object instances at a particular time (ensemble averages).

Management Communications standards

It specifies managed objects and attributes that indicate which objects and attributes are to be considered for inclusion in summary reports, scheduling of observations upon these objects/attributes, scheduling of summary reports, lists of observations to be included in summary reports, and summarization algorithms to be used. It specifies services for reporting such summaries and for initiating/terminating summarization activities.

Confidence & Diagnostic Test Classes (ISO WD 10164-y; CCITT X.737). This standard allows activation of manager-specifiable diagnostic tests (connectivity, data integrity, protocol integrity, data saturation, connection saturation, response time, loopback, function, etc.) on managed objects. These test classes permit determination of the quality of services being provided by the system being managed.

Other SMF standards that are still in their early development include the Time Management Function (ISO WD 10164-z, IS in '93; CCITT X.743), Software Management Function (10164-q, IS in '94; CCITT X.744), and Performance Management Function (CCITT X.746).

Common Management Information Service (ISO IS 9595, including IS Amendments 1,2, and CD Amendment 4 [scheduled for IS in '92]; CCITT X.710) and *Common Management Information Protocol* (ISO IS 9596, including IS Amendments 1 and 2; CCITT X.711). These standards provide the building blocks that define operations needed to perform the System Management Functions. The former, CMIS, defines services invokable by management processes when they communicate remotely. The latter, CMIP, defines the protocol to provide these services, as well as the mapping of this protocol onto the OSI remote operations service. CMIP is a powerful protocol that allows several operations to be performed in few commands, thereby minimizing management communication traffic.

CMIS services are based on the management information model. They include initialization service, information transfer service, and synchronization and linked reply services.

The initialization service supports the exchange of information about managed objects, management functions, and CMIS services supported by each end of a management association established by ACSE.

Information transfer services provide management operations services that (a) create and delete instances of managed objects (the CMIP CREATE and DELETE operations provide these services), (b) retrieve attributes of managed objects (provided by the CMIP GET operation), (c) cancel previously requested retrieval services (CMIP CANCELGET), (d) modify attributes of managed objects (CMIP SET), and (e) perform actions that may be specific to a managed object, such as initiate a diagnostic self test (CMIP ACTION). Also included are notification services that transfer events emitted by managed objects (CMIP EVENT-REPORT).

Synchronization service provides best effort or atomic synchronization of an operation across multiple objects. Linked reply service provides for multiple replies to be returned as a result of a single request and to be linked with the appropriate single specific request.

Some CMIS services have “modes”: confirmed and unconfirmed. The former requires remote management process to send a response to indicate receipt and success or failure of the operation requested. The latter doesn’t use responses.

continued on next page

OSI Systems Management (*continued*)

CMIS also describes operations that can be performed upon multiple objects, i.e., some single service requests may be concurrently applied across many managed objects. In particular, CMIS provides capabilities of selecting which objects are to be operated upon through a process of "scoping" and "filtering." Scoping identifies the pool of managed objects that are candidates to be operated upon, whereas filtering specifies which objects from the "scoped" pool will actually be operated upon. Scoping identifies candidate objects by specifying the root of a subtree in the containment hierarchy as well as the number of levels down from this root in the subtree that are to be screened (filtered). Filtering applies tests to each candidate (scoped) object to extract the particular subset that match filter criteria.

Protocol Implementation Conformance Statement (CD Amendment 5 [scheduled for IS in '92] of ISO IS 9596; CCITT X.712). This standard provides tables to be used by vendors to indicate implementation specifics, such as (a) what aspects of OSI management communication protocols were implemented, e.g., which protocol version, roles (manager, agent or both), modes and functional units, (b) application contexts and abstract syntaxes supported, and (c) limitations or ranges on values, etc.

From standards to products

Just like with other international standards, obtaining broad consensus among numerous constituencies has made many OSI management standards rife with potentially incompatible options, gaps and ambiguities. This can lead to incompatible implementations. Chappell [4] describes the commercialization process to overcome this hurdle. It fosters development of multi vendor products competitively available in any country and interoperable with comparable products from any other world region. Milestones are regionally harmonized implementation agreements profiles, called *International Standard Profiles* (ISPs), that clarify ambiguities, fill gaps and specify interoperable subsets of OSI standards. ISPs themselves become ISO standards.

OSI management ISPs are being completed in parallel with activities needed to complete remaining OSI management standards. Constituencies collaborating in ISP harmonization include the three *Regional OSI Workshops*, in which users heavily participate, as well as major vendor consortia such as the OSI/Network Management Forum, X-Open and the Japanese INTAP group. Reference [6] are North America's regional agreements serving as inputs to this international harmonization.

Procurement and testing

Other factors assisting convergence to interoperable products include recent discussions among countries to coordinate their individual network management procurement mandates. Compatible US and UK government procurement directions, based on harmonized implementation agreements, are likely by the end of 1991. Other European governments will follow suit shortly thereafter. Furthermore, to assure product conformance and interoperability, testing organizations such as the Corporation for Open Systems will begin to roll out testing products and services for these harmonized agreements also at the end of 1991.

The emerging combined pull of government and user demand, with the push of product availability, testing and interoperability, as well as public catalogues and dictionaries of managed objects, will likely be a major impetus to foster rapid fielding of all OSI technology in general.

References

- [1] *ConneXions*, Special Issue: Network Management and Network Security, Volume 4, No. 8, August 1990.
- [2] *ConneXions*, Special Issue: Network Management, Volume 3, No. 3, March 1989.
- [3] Rose, Marshall T., "Components of OSI: The Application Layer Structure," *ConneXions*, Volume 4, No. 1, January 1990.
- [4] Chappell, David, "Components of OSI: A Taxonomy of the players," *ConneXions*, Volume 3, No. 12, December 1989.
- [5] LaBarre, Lee, "Management By Exception: OSI Event Generation, Reporting, and Logging," Invited Paper in Proceedings of the *Second International Symposium on Integrated Network Management*, Washington, D.C., North-Holland Publisher, April 1991.
- [6] Stable Implementation Agreements for OSI Protocols: Part 18—Network Management, Version 4, Edition 1, NIST Special Publication 500-183, December 1990; Working Implementation Agreements for OSI Protocols: Part 18—Network Management, December 1990, NIST IR document in publication.
- [7] Benford, Steve, "Components of OSI: The Directory Service," *ConneXions*, Volume 3, No. 6, June 1989.
- [8] Marshall T. Rose, "The Simple Book—An Introduction to Management of TCP/IP-based internets," Prentice-Hall, 1990, ISBN 0-13-812611-9.

PAUL J. BRUSIL is a Principal Scientist in The MITRE Corporation's Distributed Processing Systems Division, Network Management Specialty Group and OSI Specialty Group. He is the organizer, convener and first chair of OIW's Network Management SIG and IFIP's *First International Symposium on Integrated Network Management*. He is the OIW's OSI User Representative to North America's four man delegation to the Regional Workshop Coordinating Committee working to eliminate regional incompatibilities in OSI implementations and thereby to bring globally interoperable OSI products to market as quickly as possible. Paul received his Ph.D from Harvard's Division of Engineering and Medical School in 1973.

Components of OSI

Ed.: This article is one of many in a long-running series called *Components of OSI*. Back issues are available for \$15 each. Articles to date in this series include:

ISDN	April	1989
X.400 Message Handling System	May	1989
X.500 Directory Services	June	1989
The Transport Layer	July	1989
Routing overview	August	1989
IS-IS Intra-Domain Routing	August	1989
ES-IS Routing	August	1989
The Session Service	September	1989
CLNP	October	1989
The Presentation Layer	November	1989
A taxonomy of the players	December	1989
The Application Layer Structure	January	1990
FTAM	April	1990
The Security Architecture	August	1990
Group Communication	September	1990
X.25	December	1990
The Virtual Terminal ASE	January	1991
Systems Management	April	1991

**An SNMP Stereo System:
Musical Networks, Australian Style**

by Simon Hackett

At INTEROP 90, I had the pleasure of demonstrating a device which brings something of a new twist to the use of the *Simple Network Management Protocol* (SNMP) on a TCP/IP network.

SNMP

SNMP is designed to allow vendor-independent monitoring and management of objects on a TCP/IP network. It is conventionally used to monitor and manage devices such as IP routers and bridges, and compute hosts on an IP network. It is in essence a protocol which allows an IP node to provide access, over the network, to a tree-structure of variables. Variables can be read-only or read/write. The ability to write values into variables is controlled by a fairly rudimentary scheme of "community" strings—essentially a plain-text password.

There is a standard tree of variables, a so-called *Management Information Base* (MIB), which each vendor of SNMP products should support. These allow the retrieval of information about the IP node—information such as IP network numbers, routing tables, and packet input, output and error counts.

An unusual application

New tree structures (MIBs) can be defined to suit other applications. To prove the point, I demonstrated an unusual application—a networked stereo system, using SNMP protocols to provide full monitoring and control of the system. The stereo system we demonstrated was a Pioneer Tuner/Amp (with a cute motor-driven volume knob), a Pioneer PD-M910 six-disc CD player, and a set of Klipsch speakers. We operated the stereo system for the interest of visitors to the TGV booth, using X-windows front end software plus some simple command-line based SNMP tools.

Hardware

Connecting this system to the network was a little "magic box." Based on a Motorola 68000 processor, this device performs real-time monitoring and control of the stereo system, and also runs an IP kernel, with ICMP, UDP, and an SNMP agent which implements our audio-visual MIB. It talks IP protocols to the world using serial line IP (SLIP). Interfacing the to stereo system is achieved using two connections. The device listens to the signal coming from the the "Digital Output" jack on the CD player, and generates a control signal which connects to the Pioneer-standard "remote control input" jack on the CD player. The "remote control output" on the CD player is daisy-chained to "remote control input" on the tuner/amplifier, so the device can control both units. (See Figure 1).

Software

The software in the box was written almost entirely in C, using Sun workstations. Much of it was written during a period of intense activity at TGV Inc in Santa Cruz, California during August of 1990, when I worked with several others to finish the software in the device, adding the IP and SNMP support to the control software I had previously written.

The control software uses some careful timing to imitate the remote control signals used by Pioneer's stereo components. Thus, by being plugged into the "remote control in" jack on the CD player, any CD player or Tuner/Amp function can be initiated by the controller. The digital output signal from the CD player contains both the digital audio samples for the music being played, and a stream of status information.

This information is decoded in real-time using interrupt driven routines to provide the system with continuous monitoring of the position of the CD player in any music selection played. The "table of contents" information from each CD is also read using this interface.

Using the system

The device implements a 100 entry play queue for the CD player, much like a jukebox. It processes this queue, playing selections, and allows any connected network node to monitor this queue, add selections to it, and do other monitoring and control of the stereo system (tuner/CD selection, amplifier volume adjustment, fade down, fade up, etc). SNMP operations allow full management of the queue, including retrieval of queue elements and insertion, deletion and replacement of entries anywhere in the queue.

For INTEROP, we wrote some demonstration applications to show off the unit. A set of X11 tools show the status of the system (including the currently playing disc number, the playing time of the current selection in minutes and seconds, the amplifier volume, tuner status, etc). Recognizable buttons are provided in the window to provide the play, pause, stop, eject and other required functions. (See Figure 2).

Another window displays a set of disc titles from a database of available discs. Clicking on a disc title pops up a window listing all the tracks on that disc, and also all the tracks on the other discs in that "six pack." Clicking on track titles in the window causes the appropriate SNMP requests over the network to the controller, which appends the selection to the play queue.

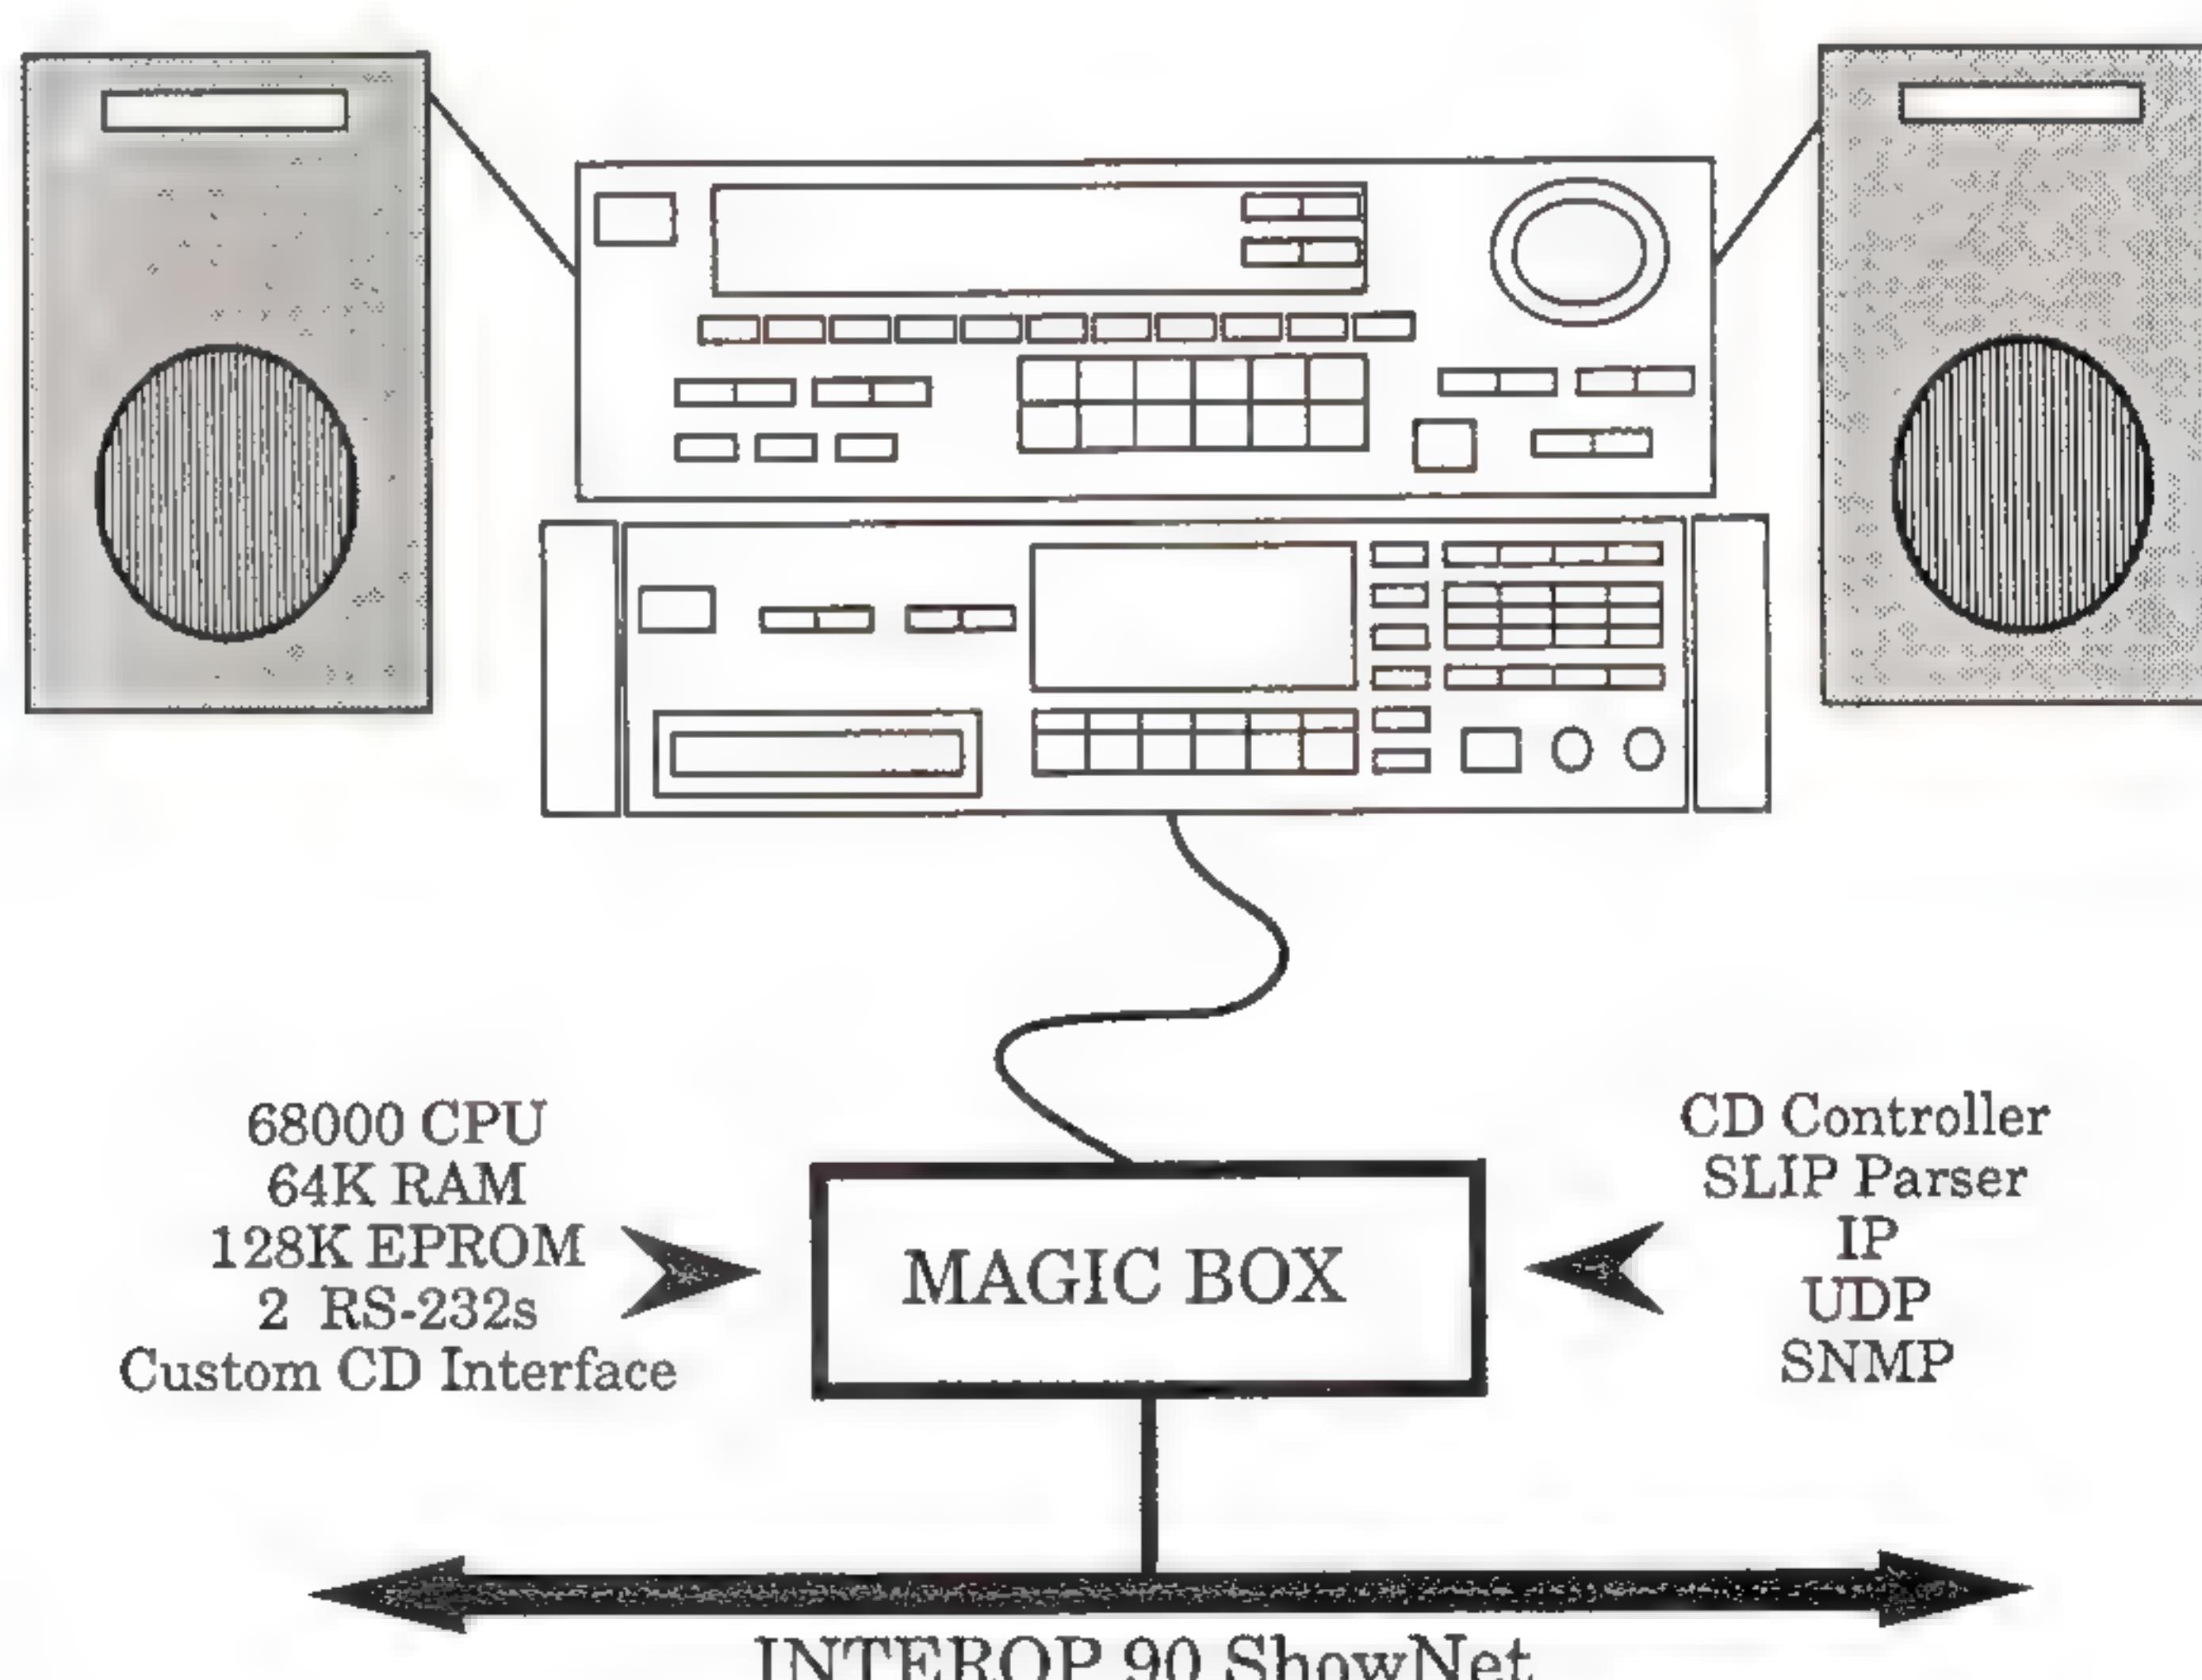


Figure 1: Hardware configuration

SNMP control issues

With Marshall Rose of PSI Inc and Stuart Vance of TGV Inc, I developed an "audio-visual" MIB, providing a tree structure which describes the functions and information available from a range of stereo components. This MIB is quite substantial, since it describes all the functions of the six disc CD player, and most functions of the Tuner/amp unit. It contains sections for most other sorts of audio/visual components, which we will "fill out" as needed.

Because the SNMP protocol is general, any software capable of issuing SNMP *set* and *get* requests is capable of accessing the stereo system, by using our custom MIB to find the identity of the required variables. This includes line-based SNMP libraries available from several sources, and SNMP network management station software.

continued on next page

An SNMP Stereo System (*continued*)

Control and monitoring of the stereo is effected with sub-trees of SNMP variables describing each function. For instance, one sub-tree contains current status information for the CD player (current disc, current track, time into track, time into disc etc). This sub-tree can be “walked” using SNMP’s *powerful get-next* operator, to display the player’s status. The values returned are real-time information from the CD player.

Another sub-tree contains information describing each disc in the six pack. If less than six discs are loaded, the sub-tree simply shrinks as appropriate. For each disc, a further sub-tree gives an identity number for that disc, the number of tracks on the disc, and the duration of each track.

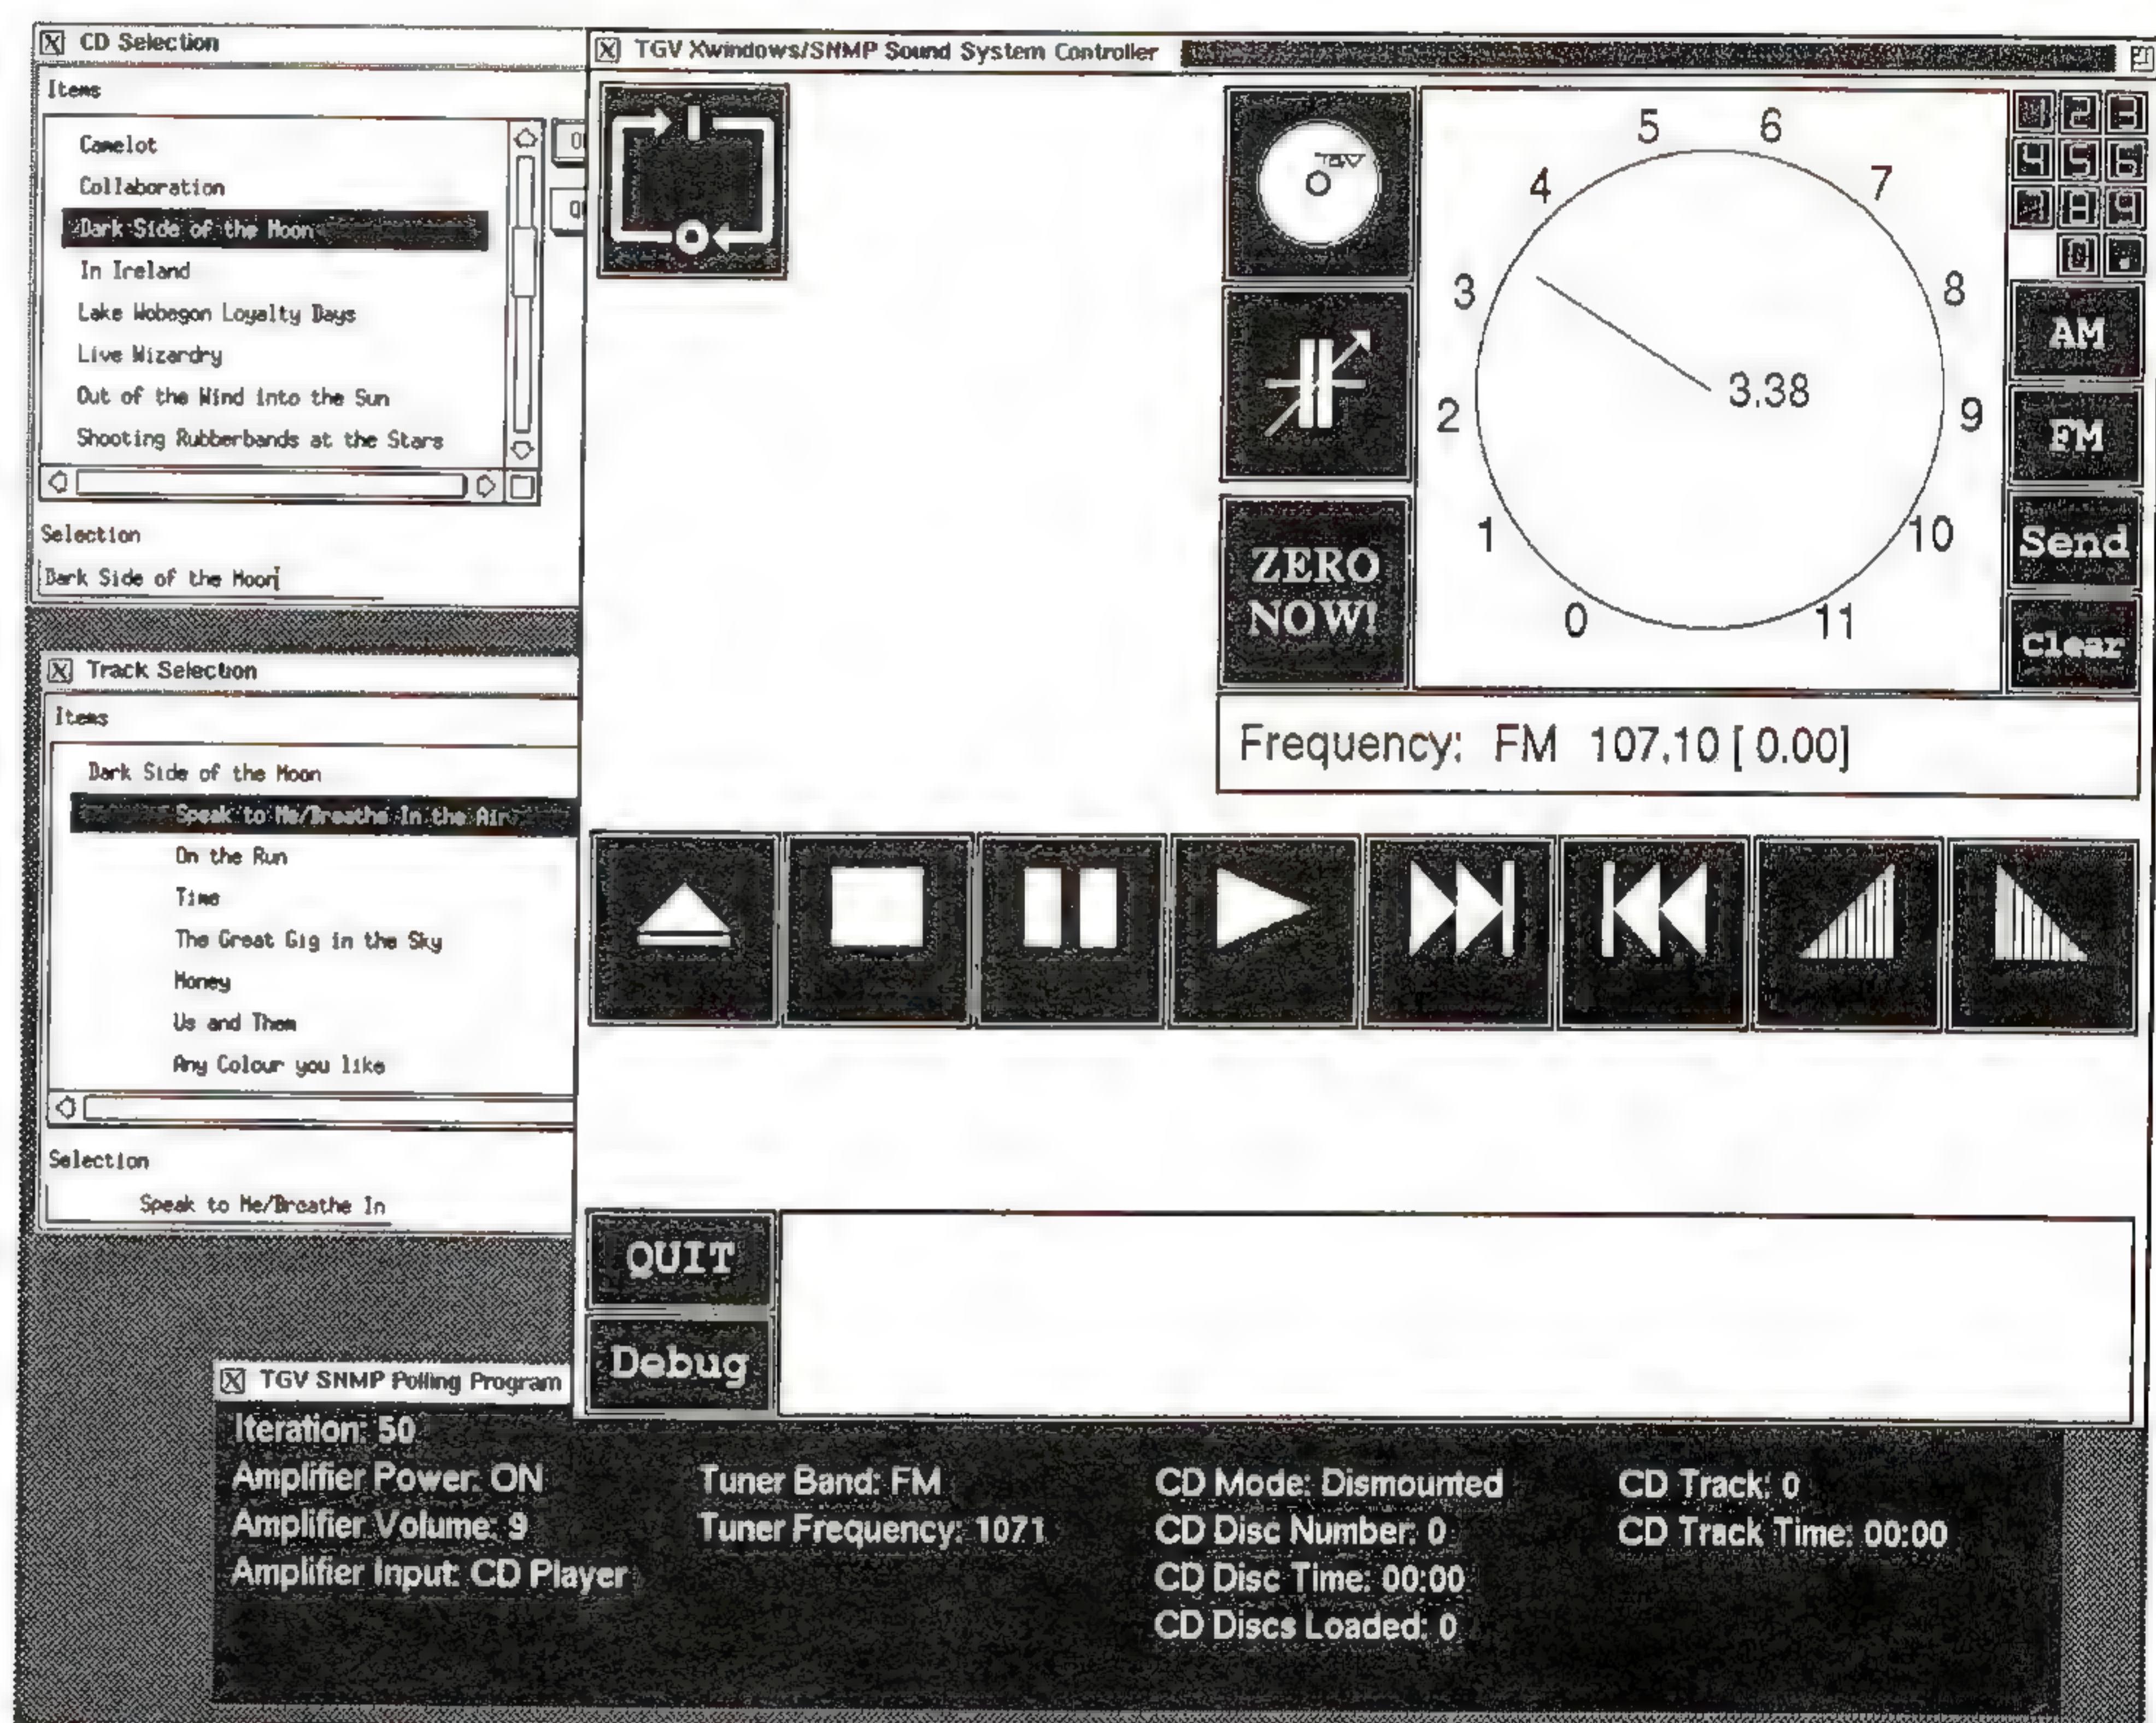


Figure 2: The X Tools CD player control window

Yet another sub-tree is responsible for the queue of tracks waiting to be played. For each entry in the queue, a part of this tree describes the attributes of that selection. These are: disc ID, track number (“0” implies the whole disc), starting and ending times, plus track repeat and requeue counts. Each entry in the queue is assigned a unique 32 bit ID by the software for later reference, and this ID also appears as an SNMP variable in each queue entry. Walking this sub-tree provides a list of the tracks waiting to be played.

To instruct the system to queue a selection, a set of SNMP variables is written in as a “request block,” and then a function code is written into a variable to tell the system to queue the entry. Functions have been implemented to allow queue entry insertion, deletion, replacement, etc. Other system settings are accessible via SNMP. For instance, one variable is the amplifier volume setting. This can be changed with an SNMP *set* request to change the volume, and read using an SNMP *get* to check the current volume level.

Wide area access

Since the controller is a functional IP node, it can be accessed anywhere on the internet it is connected to. Indeed, I have gained some amusement from altering the volume setting on a system at TGV in Santa Cruz, CA, from Adelaide, South Australia over the Internet. This is quite a good computer-hackers' party trick.

Conclusion and futures

At INTEROP 90, we demonstrated that SNMP can control anything you want to control. Indeed, the backup controller device was used to operate the *second* Internet Toaster to make an appearance at INTEROP 90; the first Internet Toaster was demonstrated at INTEROP by John Romkey. We also found out that the majority of networking people at INTEROP 90 seemed to prefer listening to Pink Floyd and Monty Python (almost to the exclusion of anything else, it seemed...)

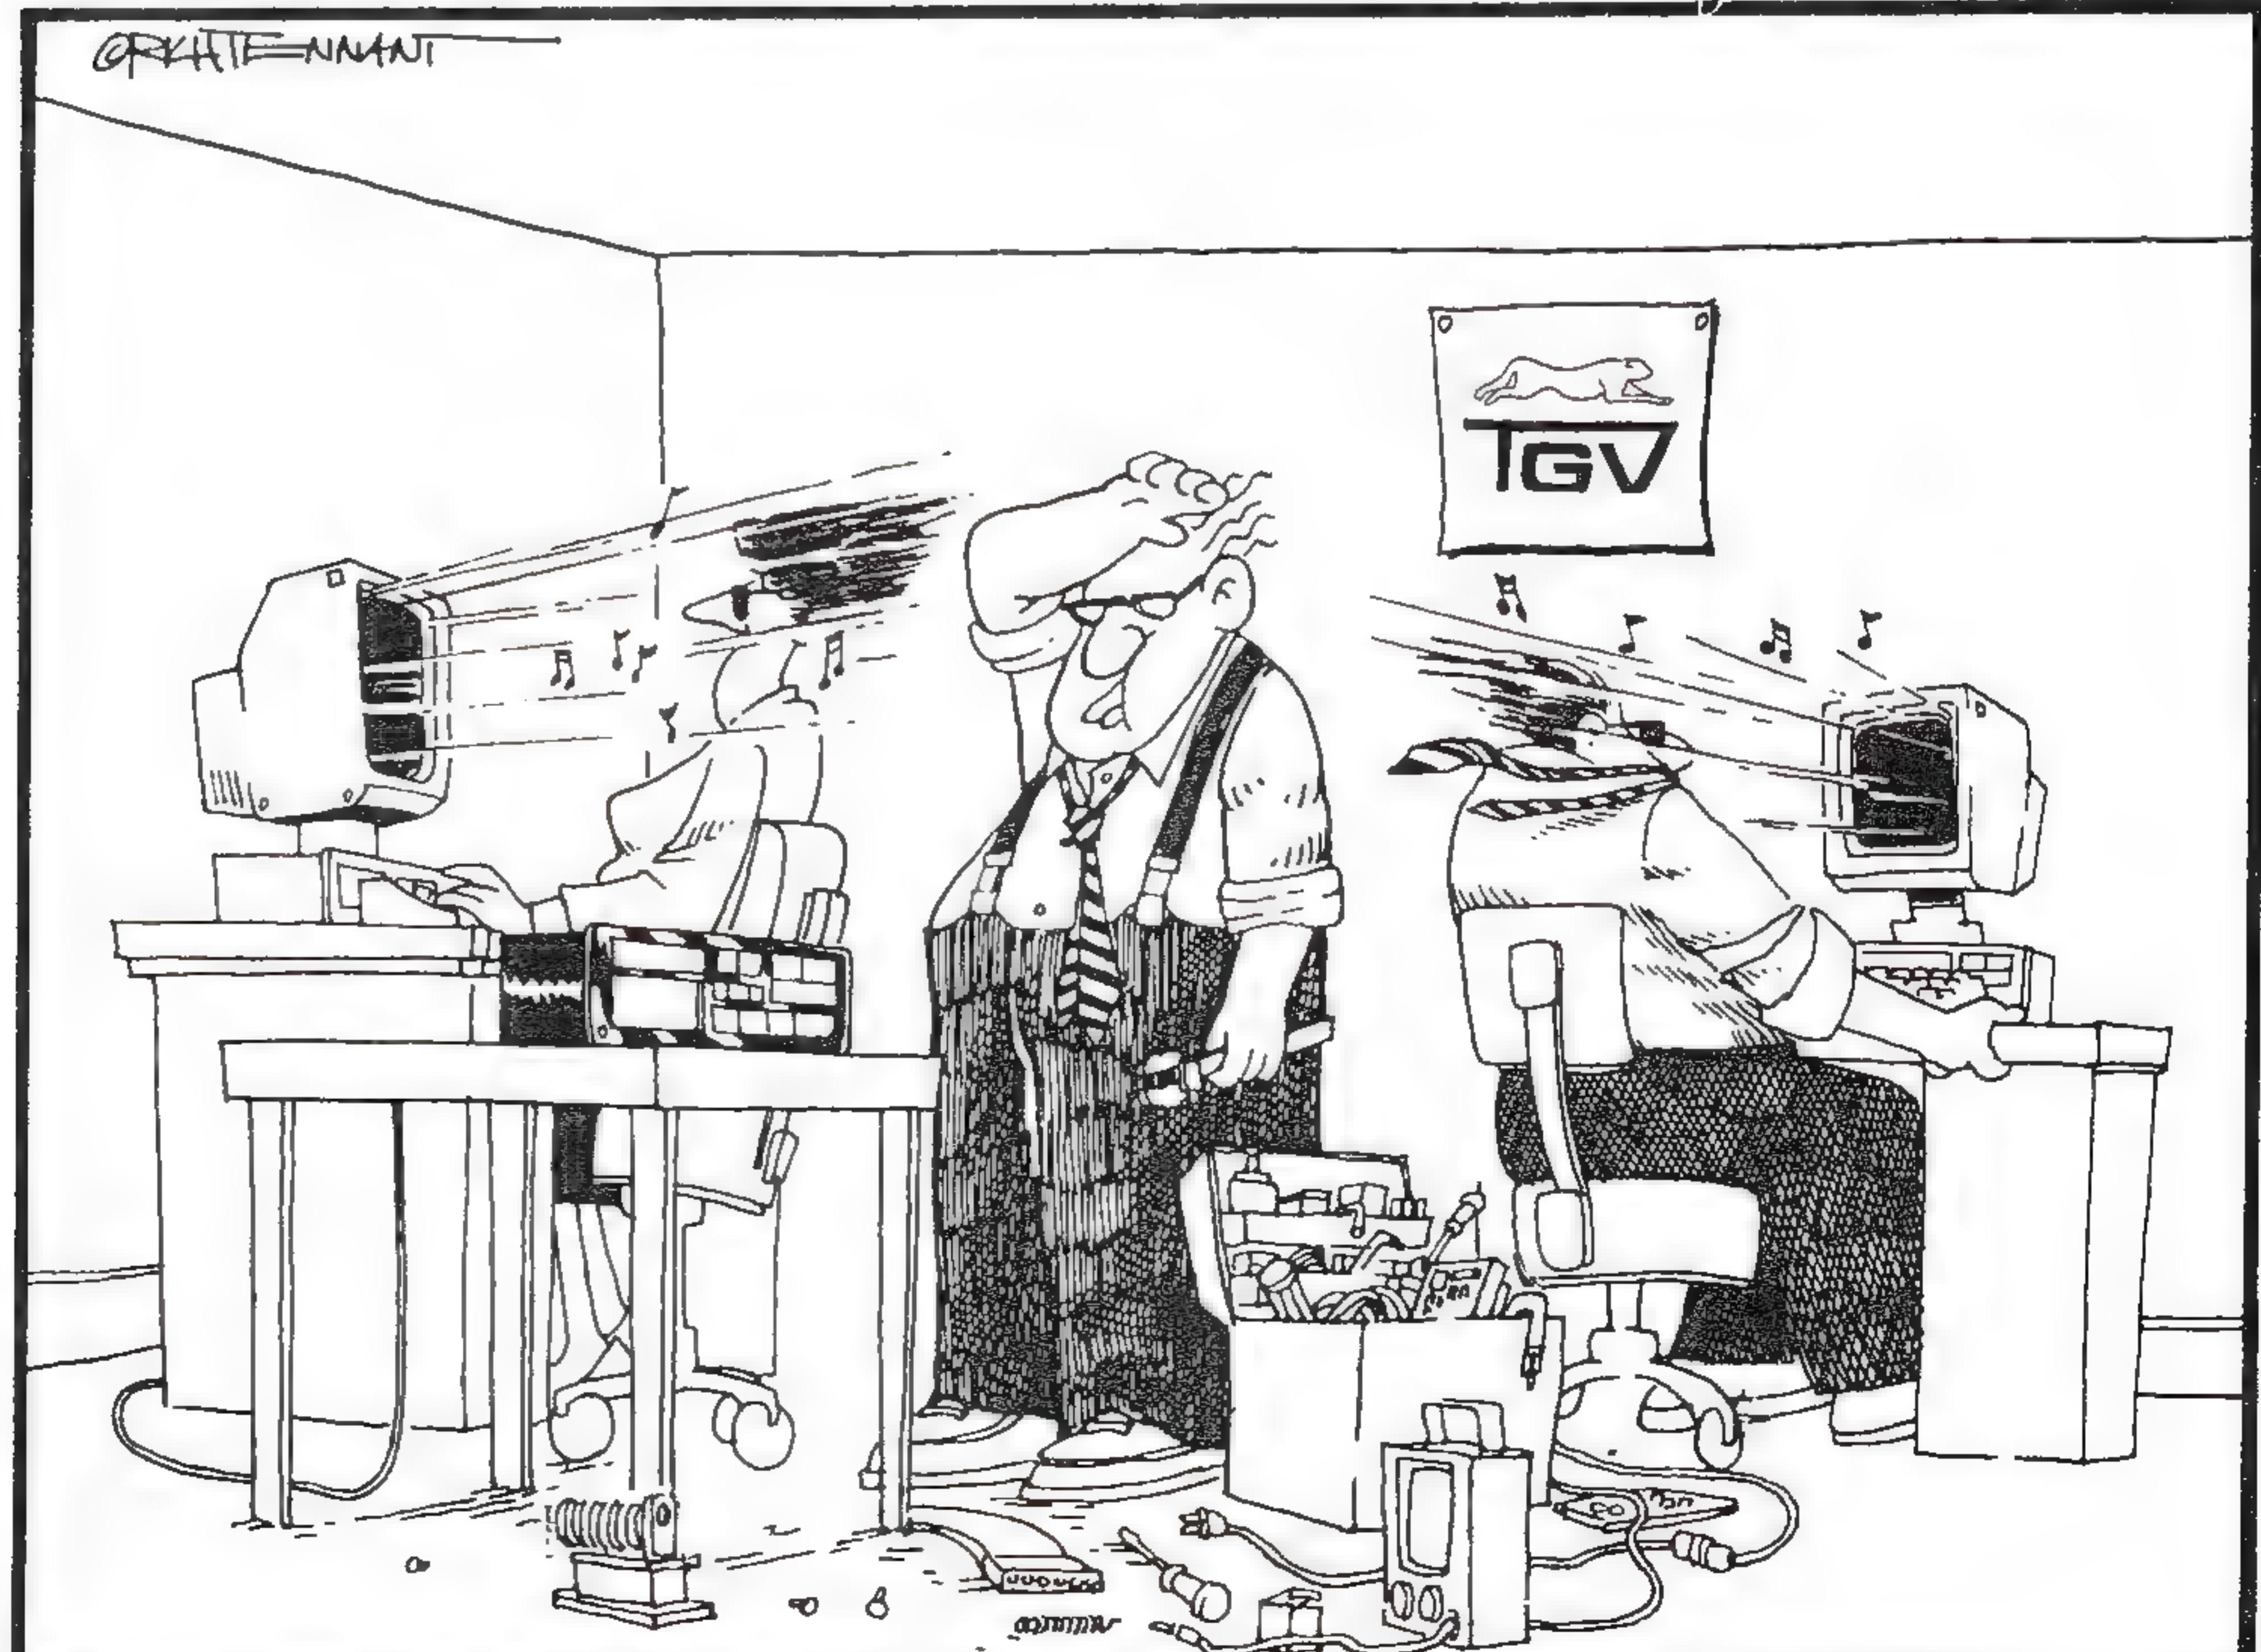
I hope to demonstrate the next generation of this device at INTEROP 91. It will be using Ethernet instead of a serial line for the IP connection, and instead of just monitoring the stereo system, I intend to push the digital audio data from the CD player down the network and pop it out of another black box somewhere else on the show floor! I should also have some other interesting applications of this technology ready to demonstrate at the show.

Acknowledgements

This enterprise would not have happened without the generosity and encouragement of TGV, Inc in Santa Cruz, California. I need to thank Ken Adelman and David Kashtan of TGV for their support (including flying me over to make this happen), John McMahon of TGV for writing the X-windows code. John Romkey and Karl Auerbach of Epilogue Technology assisted us and kindly allowed us to use their SNMP agent code as the core of our SNMP implementation. Most importantly, I am very grateful to Stuart Vance of TGV for his friendship and encouragement, and for his contribution of a great deal of his time and effort to help make this happen.

The 5th Wave

By Rich Tennant



SIMON HACKETT is an Australian consultant who works in the fields of networking, communications and machine control. He maintains the South Australian hub of the Australian Academic and Research Network (AARNet), the Australian equivalent of the Internet. Some of his hobbies include flying gliders, collecting old workstations and doing unconventional things with stereo equipment. He can be reached as:

simon@itd.adelaide.edu.au.

IN A DISPLAY OF PERVERSE BRILLIANCE, SIMON THE REPAIRMAN MISTAKES A COMPACT DISK PLAYER FOR A WORKSTATION SYSTEM UNIT, BUT MANAGES TO TIE IT INTO THE NETWORK ANYWAY.

**Convergence of European and North American
Research and Academic Networking**
by
**Steven N. Goldstein, National Science Foundation
and**
**Christian Michau, Centre National de la Recherche
Scientifique**

Introduction

This article reviews highlights in North American and European research and academic networking during 1989–90 and interprets these developments as way stations on convergent paths. Present developments and future trends are discussed. We suggest that research and academic networking developments in Europe and North America are on convergent paths. Convergence is needed to realize the fullest potential of networking to support international research and academic collaboration.

Our approach is to highlight developments in the past year or so which demonstrate a convergence between the computer networking and information processing environments in Europe and North America. The authors are optimists who prefer to see the glass as half full rather than half empty. Those who see emptiness might point to the apparent schism between OSI protocol implementation in Europe and the widespread use of the DARPA protocol suite in the US. But, closer examination will reveal developments which indicate that these seemingly disparate approaches are actually on convergent paths.

The list of highlights is based on recollections of the authors. No precedence is implied by the order of listing. Further, while the authors have attempted to report the events with fidelity, they take full responsibility for their interpretations and any resulting misrepresentations of fact.

DFN implements WIN

In Germany, the *Deutsche Forschungsnetz* (DFN) Verein has implemented *Wissenschaftsnetz* (WIN), or Science Network. 180 sites (universities and research institutions) have been connected by means of 9.6 kbps and 64 kbps X.25 links. 2Mbps service is scheduled to start in late 1990. DFN will implement an IP gateway in WIN for access to the Internet.

High-speed networking

Several high speed optical fiber based networks of metropolitan-to-regional extent have been implemented in key academic and industrial research areas of France. Interconnections with a high speed national backbone are on the near horizon. The networks are multi-protocol. There is a large IP community in France, and OSI coexistence is underway: OSI applications are in use, and CLNS is imminent.

Italy has extended a 2Mbps link to CERN, and has begun the implementation of its 2Mbps national backbone under the leadership of the Research Network Harmonization Group (GARR). The Italian network will be multi-protocol. Similarly, both the UK and Germany provide 2Mbps service either now (UK), or by the end of 1990 (Germany).

SWITCH

Switzerland started the implementation of *SWITCH*, its national network for research and education, 1987. In 1989 and 1990, the services *SWITCHmail*, an X.400 national mail backbone with gateways to all major networks, and *SWITCHlan*, a 2Mbps multi-protocol backbone supporting DoD IP, ISO CLNS and DECnet, were completed.

IXI	<i>Réseaux Associés pour la Recherche Européenne</i> (the European Organization of Research Networks, RARE), as Project Manager, and the Dutch PTT, as the network service provider, implemented the International (pan-European) X.25 Infrastructure (<i>IXI</i>) Pilot Service in April, 1990. The IXI network now connects the national academic and research networks of European countries as well as some of the packet-switched public data networks. Initially, connections are established with 64kbps to two main nodes (Amsterdam and Berne) interconnected with a double 64kbps backbone. A study of a follow-on pan-European backbone providing up to 2Mbps access points is under preparation, and a multi-protocol service is under consideration.
Emergence of RIPE	<i>Réseaux IP Européens</i> (European IP Networks, RIPE) emerged but a year ago as a confederation of production IP networks in Europe. RIPE members work within the framework of RARE and have gained acceptance for their message that there are communities of users throughout Europe who rely on IP networks to accomplish their research and academic pursuits and that coordination among them is both needed and is not inconsistent with an orderly transition to OSI protocols. As a result of the mutual understanding promoted by the leadership of both RIPE and RARE, RIPE is recognized as a group within RARE for the coordination of IP administration in Europe. RIPE has produced Terms of Reference, has set up working groups, has begun the collection and archiving of traffic statistics and network configuration data, and RIPE is designing a European IP Network Information Center (NIC). RIPE is also working on the design of a backbone to interconnect European IP networks.
NSFNET at T3 speeds	In June, the National Science Foundation (NSF) announced the addition of three new nodes to its backbone as well as the initial phase in the installation of a T3 national backbone in cooperation with the NSFNET partners, Merit, IBM and MCI.
CLNP on backbone	The NSFNET Partnership demonstrated the operation of CLNP in NSFNET routers at the INTEROP exposition in California in October, 1989. NSFNET plans to have CLNP installed in the operational backbone in 1990.
ESnet	In early 1990, the Department of Energy implemented its consolidated <i>Energy Sciences Network</i> (ESnet) national T1 backbone. ESnet, DARPA's Terrestrial Wideband Net, TWB, below), the NASA Science Network, and NSFNET have been interconnected on the east and west coasts of the US at <i>Federal Internet EXchanges</i> (FIXes) engineered by the <i>Federal Engineering Planning Group</i> (FEPG). This facilitates rational management of routing among the national network backbones.
Terrestrial Wideband Net	The Defense Advanced Projects Agency (DARPA) switched over from satellite to terrestrial service in its Wideband Net. The TWB's STream (ST) protocol implementation permits bandwidth reservation for such activities as multimedia video conferencing. ST technology will be employed in the US-UK "Fat Pipe."
Federal Networking Council	In the Spring of 1990, the ad hoc US Federal Research Internet Coordinating Committee (FRICC) was replaced with a formal entity, the <i>Federal Networking Council</i> (FNC). The FNC membership extends beyond the five-or-so member agencies of the FRICC to include more than a dozen agencies which coordinate, at a planning and policy level, the evolution of the Federal Research Internet into the <i>National Research and Education Network</i> (NREN).

European & North American Networking (*continued*)

The FNC has established working groups to carry out the detailed operational and developmental activities. Current working groups are the Engineering and Operations Working Group, the Research Working Group and the Security Working Group. Close ties have been established between the FNC and RARE in order to promote active cooperation for intercontinental networking.

CA*net

Canada began implementation of its national research and academic network, *CA*net*, in the summer of 1990. The national backbone will eventually be a two-adjacent-ring structure connecting the provincial-regional networks in Canada. *CA*net* and NSFNET both employ IBM switching technology, and the interconnection of the two networks will be directly between three *CA*net* nodes and three NSFNET nodes (Vancouver–Seattle at 56kbps, Toronto–Ithaca at 168kbps, and Montreal–Princeton at 112kbps). Like the NSFNET, *CA*net* anticipates operating with OSI protocols as they become available. Initially, *CA*net* will use IP, and the Canadian BITNET entity, will use *CA*net* to relay BITNET-over-TCP/IP ("BITNET II," also known as VMNET). Canadian HEPNET will also use *CA*net* with DECnet encapsulated in IP.

CONACYT-RAM

The Science and Technology Council (CONACYT) has begun to implement a satellite-based Mexican Academic Network (RAM). It will connect with the NSFNET at Boulder, Colorado. Eventually, CONACYT-RAM will include the existing satellite-based networks of the Mexican National University System (Red UNAM) and Monterrey Institute of Technology (ITESM).

Multi-agency, Bilateral Fat Pipes

Realizing that the [growing] proliferation of network connections (links) across the Atlantic Ocean is not only wasteful of resources, but that it could complicate routing within and among networks, government and private interests in North America and Europe have collaborated to commission jointly funded and shared circuits between the two continents.

UK-US

The first such agreement was between DARPA, the National Aeronautics and Space Administration (NASA) and NSF in the US and the Ministry of Defence (MoD) and the Joint Network Team (JNT) in the UK. Initially, the circuit will support the equivalent of eight 64kbps channels. 128kbps will be multiplexed for the support of NASA missions involving co-investigators at British universities. The remaining 384kbps will be infrastructural (general purpose). NSF-JANET and DARPA-MoD will share at least 128kbps of the infrastructure channel, or the full 384kbps when no scheduled experimental applications such as video conferencing are running. The US-UK "Fat Pipe" will enter testing in July, 1990.

FRG-US

A similar arrangement was struck among the Department of Energy (DoE), NASA and NSF in the US and DFN in Germany. In this case, the US agencies will share the costs of 128kbps with DFN for connectivity between WIN and the US Internet. The FRG-US "Fat Pipe" is scheduled for late-1990 implementation.

Trans-Atlantic T1

As part of its European Academic Supercomputer Initiative (EASI), IBM Corporation has implemented a pan-European network, EASInet to connect IBM supercomputer sites in Europe. In a related move, IBM funded the first research and academic trans-Atlantic T1 IP link, EASIGate between the EASInet hub at CERN and the NSFNET node at the NSF-sponsored IBM supercomputer site at Cornell University in New York State.

The link started operation in March, 1990. The EASInet Project Committee (EPC) has been convened by IBM as the forum for discussing EASInet and EASIgate, and in particular the conditions of access for non-EASI sites. IBM (as sponsor), GMD (as network management contractor), and CERN (as EASIgate host site) have permanent membership. Four other members (initially GSI, CINECA, CEA, SARA) represent EASI sites as a whole.

X.400 demo project

In 1989, NASA funded the University of Wisconsin Computer Science Department to introduce X.400 experimentally in selected Internet sites. The plan is to operate X.400 implementations over TCP/IP communications layers initially and to extend to CLNS when it is available end-to-end in the Internet. NSF was among the initial installations of the Wisconsin X.400 package. Larry Landweber and RARE WG1 reached agreement for connecting the Wisconsin X.400 hub to RARE MHS networks, initially through the Norwegian well-known entry point (WEP). Successful connections were made in early 1990.

Experimental X.500

Performance Systems International, Incorporated (PSI) and NYSERNet undertook an experimental X.500 Directory Service Pilot Project in 1989 under the direction of Marshall Rose. Many US agencies (including NSF) and corporations have entered their directory information in the PSI X.500 data base which is accessible over the Internet, including an X-Windows graphical interface.

Commercial IP ventures

The last year has seen the launching of several IP networking ventures serving the commercial sector as well as the academic and research sectors. The Finnish PTT in Europe and *PSInet* and *Altinet* in the US are among the pioneers in commercial IP networking.

Networking with Eastern Europe and USSR

As the Berlin Wall fell, so did many barriers to academic and research networking between West and East Europe. Hungary and Czechoslovakia have joined RARE. Both Poland and the Soviet Union have been invited into RARE membership, but neither country has responded yet. Poland, Hungary, Czechoslovakia and the Soviet Union have been invited to join EARN. Initial EARN connections to these countries have been implemented or are underway. The principal barriers (reportedly) are the availability of communications circuits and/or their affordability in hard currencies. Also, EUNET links to Czechoslovakia and Hungary are operational while others are presently being discussed or implemented.

PACCOM

In the Pacific, Australia, Japan, Korea, New Zealand and Hawaii have forged a Pacific Computer Communications network consortium, *PACCOM*. *PACCOM* is based on Ethernet bridging and dual protocol (IP and DECnet) routers. *PACCOM* connects to the US Internet through a Hawaii-to-NASA Science Internet circuit.

Plans to add other Pacific Rim countries are underway. Japan has implemented several internal IP networks, notably *WIDE*, *TISN*, *ICOT* and *ISR*. Australia has implemented its Australian Academic Research Network (*AARN*) as a two-level-star multi-protocol network connecting the country's regional networks to Melbourne. Korean (*SDN/KAIST*) and New Zealand (*UNINET*) national academic and research networks also receive international connectivity through *PACCOM*.

Present developments**European & North American Networking (*continued*)**

Several activities are now underway which evidence greater cooperation among the leadership of both continents. In the applications arena, there is continuing cooperation in developing secure electronic mail among US and European contingents. Developers are hindered by national restrictions on the export of encryption capabilities needed to support authentication and privacy. But, the developers continue to talk and to cooperate with each other. There is also cooperation in the difficult area of X.400 O/R address space organization: how it [does/doesn't, should/shouldn't, etc.] map with X.500 Directory names as well as RFC 822 mapping, and the many philosophical arguments about separating routing information from naming information in any type of address. There is no obviously "right" solution (though there are abundant claims of rectitude), and many people on both sides of the Atlantic continue to grope within their own administrative environments and with their intercontinental colleagues toward achieving a pragmatic solution.

CO-CL interworking

Last year European and North American network administration and planning officials agreed to several workshops to effect cooperation. The first in the series, a RARE/FNC-sponsored Connection-Oriented/Connection-Less (CO/CL) workshop took place in July near Washington. It was small and attended by networking specialists. (A workshop on developing X.500 implementations is also planned.) Meanwhile, RARE WG4 is preparing an OSI CLNS pilot operation in cooperation with NORDUnet, CERN, SURFNET, SWITCH, and NSFNET. In Europe the pilot will use the existing RIPE and IXI infrastructures to connect the participating networks. In the NSFNET backbone, the CLNS experiment is coordinated by MERIT. Thus, the lower levels of the OSI stack appear to be receiving considerable collaborative attention.

In the engineering and planning arena, RARE has agreed to sponsor a European Engineering Planning Group to provide network engineering and operations advice in Europe (for example, engineering the high-speed multi-protocol pan-European backbone) in a manner similar to the activity of the Federal Engineering Planning Group in the US (which has engineered the interconnections among the federal agencies' national research backbones and which is helping to plan for the federal portion of the National Research and Education Network). Further, both planning groups would establish intercontinental liaison through an Intercontinental Engineering Planning Group (note that a North American Engineering Planning Group has not yet been formed but that close coordination is maintained with Canadian and Mexican developments). This structure parallels that of the North American, European and combined Coordinating Committees for Intercontinental Research Networking (CCIRN) which are composed of network administration and planning officials.

In Europe, RIPE is planning a European NIC which will carry out IP administration for Europe. At the same time, the Internet Activities Board (IAB), which has been responsible for guiding the development of the DARPA protocol stack, is readying recommendations which would help to make Internet administration more international. These activities are thoroughly complementary. In addition to the operational pragmatism involved, the global distribution of IP administration authority acknowledges the important role of the international community in authorizing access to and membership in the Internet.

Network Management

There has been less cooperative work on network management. This may in part reflect the very different outlooks of European and North American networking communities: in countries where another authority (e.g., the PTT) provides "the network," the administrators and the developers can focus their attention on the applications, whereas in the connectionless world, administrators and developers concern themselves more with building and operating the network and less with the applications. Ask an IP network specialist to describe his/her network, and one is likely to receive a description of the topology of routers and circuits; pose the question to users of X.25 networks, and one is likely to learn of the topology of Message Transfer Agents and Directory System Agents. However, the focus on joint operation of trans-Atlantic circuits which connect dissimilar networks could well serve as the vehicle to direct more joint effort toward management issues. In this regard, the participation of RIPE may facilitate matters, because many individuals who are active in RIPE also have a hand in the operation of connection-oriented networks. Such individuals can help to bridge network management development between the connectionless and connection-oriented domains, and, in the process, between North America and Europe.

Conclusions

Although the academic and research networking environments in Europe and North America are different in many ways, there is ample evidence of convergent paths. A key ingredient, which both transcends and nurtures the technological developments, is the mutual perception of consistent policies and policy implementations on both continents. In North America there is visible evidence that OSI protocols are taking root in a meaningful way. And Europe, reading those signals, as well as seeing the pragmatic value of using mature portions of the DARPA stack to meet current user needs, displays an official flexibility that few would have believed possible just a few months ago.

It is likely that multi-megabit, multi-protocol continental backbones will emerge in the next few years on both sides of the Atlantic Ocean. Vendors are already delivering IP/CLNP routers, and IP/DECnet multi-protocol routers have been around for several years. Both connectionless and connection-oriented networks will run concurrently in Europe, while connectionless service will dominate in North America. Indeed, the greatest challenge is likely to be interworking between connectionless and connection-oriented domains.

Gateways and dual stacks will permit protocol coexistence and will facilitate staged migration; they will no doubt facilitate cooperation among communities by offering alternatives to "religious wars" over protocol policies. As a result, the users will benefit from wider connectivity than would have been possible in single-protocol communities.

Compatibility on both sides of the Atlantic (as well as in the Pacific) will increase markets and can be expected to spur product and service development. The transitions to 45Mbps and beyond into the hundreds of Mbps regimes may be accompanied by service offerings (such as asynchronous transfer mode—ATM—switching) that will obsolete the schisms between connectionless and connection-oriented services. Many network researchers believe that wholly new networking protocols will be needed to exploit the Gigabit regime. Given the consolidating intercontinental market, it is likely that the protocols can be developed, tested, and implemented in a manner that will finish the task of convergence.

Glossary of abbreviations

European & North American Networking (*continued*)

AARN Australian Academic Research Network

Alternet commercial, US-based international IP network

CA*net Canada's research and academic network (Note that the * denotes the maple leaf, Canada's national symbol.)

CCIRN Coordinating Committee for Intercontinental Research Networking

CEA Commissariat à l'Energie Atomique (French Atomic Energy Commission. Geographically, the CEA site at Saclay is on EASInet)

CERN European Laboratory for Particle Physics (Geneva, Switzerland)

CINECA Consorzio Interuniversitario del Nord Est per il Calcolo Automatico (i.e., the regional academic computer centre in Bologna, and a major site in GARR)

CL as in CLNS or CLNP, connectionless

CLNP protocol to provide Connectionless-mode Network Service

CLNS Connectionless-mode Network Service

CNRS Centre National de la Recherche Scientifique (French National Center for Scientific Research)

CONS Connection-mode Network Service

CONACYT Consejo Nacional de Ciencia Y Tecnología (National Council of Science and Technology, Mexico)

COSINE Cooperation for Open Systems Interconnection Networking in Europe

DARPA US Defense Advanced Research Projects Agency

DECnet Digital Equipment Corporation's proprietary networking protocol suite

DFN Deutsche Forschungsnetz Verein (German Research Network Association)

DoE US Department of Energy

EASI European Academic Supercomputer Initiative

EASIgate EASI T1 connection between CERN and the NSFNET

EASInet network of EASI connections in Europe

EPC EASI Project Committee

ESnet US Department of Energy's Energy Sciences Network

ETH Eidgenössische Technische Hochschule (Swiss Federal Institute of Technology, Zürich)

EUNET European UNIX Users Network (Headquartered in Amsterdam)

EUREKA European Research Coordination Agency

FEPG US Federal Engineering Planning Group

FIX Federal Internet Exchange, an engineered and managed interconnection of the US federal agency research network backbones

FNC US Federal Networking Council

FRG Federal Republic of Germany

FRICC US Federal Research Internet Coordinating Committee

GARR Gruppo Armonizzazione Rete per la Ricerca (Research Network Harmonization Group, Italy)

GMD Gesellschaft für Mathematik und Datenverarbeitung (Institute for Mathematics and Data Processing, Germany)

GSI Gesellschaft für Schwerionenforschung (German Federal Research Institute for Heavy Ion Research)

HEPNET High Energy Physics Network (an international research network with a high proportion of DECnet service)

IAB Internet Activities Board

ICOT Japanese Fifth Generation Computer Project

ITESM Monterrey Institute of Technology (Mexico)

IXI International X.25 Infrastructure (Europe)

JANET Joint Academic NETwork (UK)

JNT Joint Network Team (UK)

KAIST Korean Advanced Institute of Science and Technology

MoD Ministry of Defence (UK)

NETNORTH Canada's BITNET association

NIC Network Information Center, often denoting the Internet's NIC currently operated by SRI, International in Menlo Park, California. Several other NICs exist.

NORDUnet International backbone connecting the national research and academic networks of the Nordic countries

NREN National Research and Education Network (US)

NSFNET National Science Foundation Network (US)

NYSERNet New York State Education and Research Network (US), one of the mid-level networks of the NSFNET

PACCOM Pacific Computer Communications networking consortium

PSInet Performance Systems International, Incorporated commercial IP network

PTT Postal, Telephone and Telegraph company, the national communications monopoly in many countries

RAM Red Academica de Mexico (Mexican Academic Network)

RFC 822 Request for Comments, # 822 (Internet electronic mail address format)

SARA Stichting Academisch Rekencentrum Amsterdam (Amsterdam Universities Computing Centre, The Netherlands)

SDN System Development Network (Korea)

STream a stream-oriented network layer protocol used for bandwidth reservation

SURFNET The Netherlands' national research and academic network

SWITCH Swiss National Network for Research and Education

TISN Tokyo International Science Network (Japan)

TWB DARPA's Terrestrial Wideband Network

UNINET New Zealand's national research and academic network (not to be confused with Norway's UNINETT research and academic network)

WEP national Well-known Entry Point in the RARE experimental MHS

WG1 Rare Working Group 1, Message Handling Systems

WG4 RARE Working Group 4, Network Operations and Management

WIN Wissenschaftsnetz, Science Network managed by DFN (Germany)

XNREN an experimental X.400 private management domain managed by the University of Wisconsin's X.400 project team

[Ed.: This article is adopted from a paper which appeared in the proceedings of the IFIP 6.5 conference on *Message Handling Systems and Application Layer Communication Protocols*, held in Zürich in October 1990. Conference Proceedings are available from North-Holland Publishing Company. Printed with permission].

The Ultimate File System: The Role of Religion in the Enterprise-Wide Network

by Carl Malamud

[Editor's Note: In our November 1990 issue (Volume 4, No. 11) we looked at the question of FTAM, FTP, and NFS as three candidates for the enterprise-wide file access mechanism. Eric Fleischman argued that NFS and FTP were poorly suited for this role and nominated FTAM as the enterprise-wide solution. In this article, Carl Malamud argues that comparing NFS and FTAM is like comparing apples and oranges].

Introduction

Mangoes or Orangutans?

The Ultimate File System

It has recently been in vogue to compare FTAM to NFS and FTP. Comparing any protocol to FTP is kind of like beating up an old lady instead of helping her across the street. Let us, for the time being, leave FTP standing on the corner waiting for an escort, and look instead at FTAM and NFS.

Comparing NFS and FTAM is like comparing mangoes to orangutans. If you pick a high enough metaphysical viewpoint, they both serve the same function—both are carbon-based life forms. With a given set of criteria, one can easily compare the two and come up with a clear winner: "The mango is more portable and is thus preferable to the orangutan."

If you take a lower perspective, differences between the mango and the orangutan start to show up. The mango certainly tastes better, but you can't take a group of kids to the zoo to watch a mango and expect them to stay entertained.

It is easy to fall into the mango trap when looking at emerging standards like FTAM. When we see an international standard emerging, it is tempting to say that it will solve all problems and that it should be used to the exclusion of other protocols. After all, we don't want duplication of effort.

The question is the universe used for the comparison. Analyzing enterprise-wide file access mechanisms to pick a winner requires some limitation of scope. Of course, one could compare Telnet to FTAM and decide that FTAM was better: even die-hard FTAM fanatics would consider such an analysis silly.

In the area more traditionally known as file access mechanisms, one can pick a small subset, such as FTP, FTAM, and the Network File System. One could extend this analysis to include things as the Carnegie Mellon's *Andrew File System* (AFS), DEC's *Data Access Protocol*, Novell's *NetWare Core Protocols*. One could (but hopefully would not) even go so far as to consider things like RJE over Bisync or custom COBOL programs over X.25 as candidates.

To ease this analysis, I would like to contribute another imaginary candidate which I'll call the *Ultimate File System* (UFS). UFS is clearly better than all other candidates: it is fast, efficient, powerful. It slices, it dices, and can be implemented quickly and easily. It has an installed base only slightly smaller than FTAM.

Since UFS is clearly better, one could do a blow by blow comparison to FTAM and NFS. There is no doubt that, given any set of criteria, one can find holes in each of them. The basic problem with choosing a single enterprise-wide solution between candidates like UFS, NFS, and FTAM is that they really do different things.

The ISO Reference Model

To put this analysis in more technical terms, it would help to refer to the *ISO Reference Model*. The ISO Reference Model has often been extended past layer 7 to add layers 8 and 9: "Financial" and "Political." I would like to take this opportunity to add yet another layer, 10: "Religious." Let's start at layer 10 and work our way down.

At the religious layer, the analysis of file access protocols attempts to pick a single candidate for remote data access. This single candidate for truth, beauty, and transparent access to remote data is then deployed throughout the enterprise-wide network.

At the political layer, good reasons for employing multiple protocols begin to show up. No matter how widespread a given standard becomes, there are always groups of people that want to do things differently.

A single enterprise-wide file transfer mechanism has all the disadvantages we've found with other forms of highly centralized management decisions: inflexibility to change and making decisions based on the lowest common denominator. The whole point of distributed systems is to allow management tools to adapt to changing environments (which, by definition, change differently in different places).

At the financial layer, there are even more compelling reasons to allow multiple file access protocols. Take FTP versus FTAM, for example. For the time being at least, one can make a strong argument that FTP is more efficient (and thus costs less) than FTAM for basic file transfer activity. In a local area environment, NFS is certainly is more efficient than either FTP or FTAM.

There are other financial considerations besides performance (with the concomitant need for less hardware, software, and bandwidth). If you need a cheap, easy solution, mature protocols have more public domain implementations, making them more widely available to the general public. Vendors tend to bundle in older protocols with their operating systems, whereas the latest trend in network access tends to require a separate purchase order number.

Layer 7

Let us leave the upper layer analysis and take the viewpoint of the applications layer. Here the relevant questions are strictly (one would hope) technical. It is clear that you can take FTAM and make it the basis of a distributed file system. You could build FTAM into ROM on diskless workstations and have network-based swapping using FTAM. There is no doubt that FTAM would be found lacking in these fields of application.

Likewise, there are things that NFS doesn't do as well. Because NFS is a means for extending the local file system, it is not as fully general as FTAM. We can refer to this lack of generality as UNIXisms, but a better analysis would be to realize that NFS is a file service and FTAM is a record service.

Data Independence

Let's begin our comparison of FTAM and NFS by looking at the question of *data independence* across machine architectures. FTAM has the concept of document types. In theory, you can define any structure for data in a file and have that data move transparently across machine architectures.

In reality, FTAM applications are limited to a few basic document types: straight uninterpreted binary, and text terminated by carriage returns at the end of the line are two common examples.

continued on next page

The Ultimate File System (*continued*)

If we look at NFS, we will see that there is no concept of document types. Is this a problem? Not if we understand that the role of NFS is to leave that interpretation to the client. Do you really want the operating system for your diskless VAXstation encoded so a Sun could read it? What would it do with it?

Why make the user of the data interpret the data? It certainly reduces server load. It promotes the ability of a client to store arbitrary data on arbitrary servers. Finally, the client can access raw blocks at speeds much better than the structured access: ideal for things like diskless nodes and other applications where the client is accessing programs and not data.

This is not meant to imply that you don't have independence of data on a network in NFS. The *External Data Representation* (XDR) allows a user to take data structures and encode them for network presentation. This RPC function of encoding data is used by many applications, such as Frame desktop publishing or the Legato network backup utility.

Should NFS be structuring files? Binary data doesn't have much interchange across machine architectures. Either the application knows how to use it or it doesn't. If you need interchange of arbitrary data structures, you're looking at an RPC question, not a distributed file system.

Let's stop and look at this point again: NFS and FTAM do different things so they have a different approach to data translation. NFS is a file server; FTAM is a structured, record-oriented virtual file system.

There are some other differences between the two. FTAM uses a connection-oriented, stateful approach, whereas NFS is stateless and uses (at least in most implementations) the UDP transport layer in TCP/IP.

Again, the protocols do different things. FTAM is inherently stateful and has a built in lock facility. NFS allows the locking mechanism to be an independent service. This doesn't mean that an operating system is going to allow multiple users to walk all over data: just because the NFS protocol is stateless doesn't mean that the server doesn't keep state information through a lock manager, a duplicate request cache, or application semaphores.

Note that making NFS stateless simplifies its role as a file service, making crash recovery and other operations significantly easier. This doesn't mean that NFS can't work over a connection-oriented transport service: the Reno tape for the new Berkeley UNIX features a TCP-based implementation of NFS that works quite well in a wide-area environment.

Can NFS and FTAM coexist in a single network? I've seen lots of DEC systems that use the *Record Management Services* (RMS) to tie together NFS, FTAM, DEC's Data Access Protocols, VAX Clusters, and DEC's Distributed File System into a single coherent view of data in a wide variety of environments. The local file access mechanism, in this case the Record Management Services is responsible for masking the different protocols from the client application.

A final comment on NFS versus FTAM is the often-mentioned bromism that "NFS does not scale well." It is true that a UDP-based transport connection for large files does not work well in a wide-area environment: hence the Reno distribution with the TCP transport.

Stateful and Stateless Protocols

Scaling Protocols

In a local or enterprise-wide network, I would argue that NFS has certainly scaled quite well. Technologies like the *automounter* and the *NIS name service* make resource discovery and automatic mounting simple. Dedicated file servers from companies like Auspex and Epoch allow terabytes of NFS-accessible data to be available to clients.

Lists of Three

It is sometimes tempting to group all applications on all networks into three categories: data access, mail, and virtual terminals. Once this taxonomy is heard enough times, it starts to make sense. Every protocol is put into one of the three categories. The next step is often to find the single best protocol within the three categories. The result is a list of the three standard protocols that make up a global solution to all our problems.

This approach is simplistic at best and can be quite destructive. Computer networks do many different things for many different populations of users. Picking a simple protocol is nice, but using it to the exclusion of complementary services just reduces the user's capability to accomplish useful work.

FTAM versus COBOL

Take FTAM. Making that the enterprise-wide solution ignores distributed file systems, distributed databases, transactions processing, and a wide variety of other data access paradigms. FTAM resembles COBOL in many respects. Given a universe of COBOL, RPG II and BASIC, standardizing on one or the other helps. But specifying COBOL over complementary languages just ossifies the organization.

Let me give you an example. I used to work in the research division of the Board of Governors of the Federal Reserve System. We were trying to provide an alternative to MVS/TSO for econometric analysis and needed to hire some programmers to port our applications to UNIX. Unfortunately, MIS had decided that the standard language skill needed to work at the Fed was a facility with COBOL. Needless to say, econometric forecasting with COBOL is a little tough. Because MIS had decided COBOL was our enterprise-wide standard, Personnel wouldn't let us advertise for FORTRAN or C skills. Instead, we had to put an ad out looking for COBOL programmers and hope that somebody listed FORTRAN or C on their resume.

Enterprise-wide solutions should not attempt to provide a lowest common denominator or form a rigid model. Simplistic solutions to difficult problems don't help anybody. This doesn't mean we don't need standards. Standards are great for interoperability. Standards should be a means to interoperability, however, not a corporate religion. A given enterprise may wish to declare a corporate religion. FTAM may be an article of faith within a corporation, but religions are often difficult for the populace to live up to. Ideals to contemplate, perhaps, but in the real world truth and beauty often lose out to whatever seems best at the time.

UFS+ ?

Let's return to the question of the Ultimate File System. I'm not going to recommend that you discard FTAM and NFS in favor of UFS, even though UFS is clearly better. Instead, I would begin implementing UFS on systems in tandem with FTAM and NFS. I might even make a UFS to FTAM gateway system to ease transition. As more and more of my users begin to use UFS, I might even make it one of several required applications on any host on my network. But I certainly wouldn't say this was my network-wide, enterprise-wide solution. After all, UFS+ will be available soon afterwards.

CARL MALAMUD considers himself a secular humanist at layer 10 of the ISO Reference Model.

CONNEXIONS
 480 San Antonio Road
 Suite 100
 Mountain View, CA 94040
 415-941-3399
 FAX: 415-949-1779

FIRST CLASS MAIL
 U.S. POSTAGE
 PAID
 SAN JOSE, CA
 PERMIT NO. 1

ADDRESS CORRECTION
 REQUESTED

CONNEXIONS

EDITOR and PUBLISHER Ole J. Jacobsen

EDITORIAL ADVISORY BOARD Dr. Vinton G. Cerf, Vice President,
 Corporation for National Research Initiatives

A. Lyman Chapin, Chief Network Architect,
 BBN Communications Corporation

Dr. David D. Clark, Senior Research Scientist,
 Massachusetts Institute of Technology

Dr. David L. Mills, Professor,
 University of Delaware

Dr. Jonathan B. Postel, Communications Division Director,
 University of Southern California, Information Sciences Institute

Subscribe to CONNEXIONS

U.S./Canada \$150. for 12 issues/year \$270. for 24 issues/two years \$360. for 36 issues/three years

International \$ 50. additional per year (**Please apply to all of the above.**)

Name _____ Title _____

Company _____

Address _____

City _____ State _____ Zip _____

Country _____ Telephone () _____

Check enclosed (in U.S. dollars made payable to CONNEXIONS).

Visa MasterCard American Express Diners Club Card # _____ Exp.Date _____

Signature _____

Please return this application with payment to: CONNEXIONS

Back issues available upon request \$15./each
 Volume discounts available upon request

480 San Antonio Road, Suite 100
 Mountain View, CA 94040 U.S.A.
 415-941-3399 FAX: 415-949-1779

CONNEXIONS